Table of Contents

Introduction 1.1
General HTTP Handling 1.2
HTTP Interface 1.3
Databases 1.4
- To-Endpoint 1.4.1
- Management 1.4.2
- Notes on Databases 1.4.3

Collections 1.5
- Creating 1.5.1
- Getting Information 1.5.2
- Modifying 1.5.3

Documents 1.6
- Basics and Terminology 1.6.1
- Working with Documents 1.6.2

Edges 1.7
- Address and ETag 1.7.1
- Working with Edges 1.7.2

General Graph 1.8
- Management 1.8.1
- Vertices 1.8.2
- Edges 1.8.3

Traversals 1.9
AQL Query Cursors 1.10
- Query Results 1.10.1
- Accessing Cursors 1.10.2

AQL Queries 1.11
AQL Query Cache 1.12
AQL User Functions Management 1.13
Simple Queries 1.14
Async Result Handling 1.15
Bulk Import / Export 1.16
- JSON Documents 1.16.1
- Headers & Values 1.16.2
- Batch Requests 1.16.3
- Exporting data 1.16.4

Indexes 1.17
- Working with Indexes 1.17.1
- Hash 1.17.2
- Skiplist 1.17.3
- Persistent 1.17.4
- Geo 1.17.5

- Fulltext 1.17.6
- Transactions 1.18
- Replication 1.19
  - Replication Dump 1.19.1
  - Replication Logger 1.19.2
  - Replication Applier 1.19.3
  - Other Replication Commands 1.19.4
- Sharding 1.20
- Monitoring 1.21
- Endpoints 1.22
- User Management 1.23
- Tasks 1.24
- Agency 1.25
- Miscellaneous functions 1.26

ArangoDB v3.1.18 HTTP API Documentation

Welcome to the ArangoDB HTTP API documentation! This documentation is for API developers. As a user or administrator of ArangoDB you should not need the information provided herein.

In general, as a user of ArangoDB you will use one of the language drivers.

General HTTP Request Handling in ArangoDB

Protocol

ArangoDB exposes its API via HTTP, making the server accessible easily with a variety of clients and tools (e.g. browsers, curl, telnet). The communication can optionally be SSL-encrypted.

ArangoDB uses the standard HTTP methods (e.g. *GET*, *POST*, *PUT*, *DELETE*) plus the *PATCH* method described in RFC 5789.

Most server APIs expect clients to send any payload data in JSON format. Details on the expected format and JSON attributes can be found in the documentation of the individual server methods.

Clients sending requests to ArangoDB must use either HTTP 1.0 or HTTP 1.1. Other HTTP versions are not supported by ArangoDB and any attempt to send a different HTTP version signature will result in the server responding with an HTTP 505 (HTTP version not supported) error.

ArangoDB will always respond to client requests with HTTP 1.1. Clients should therefore support HTTP version 1.1.

Clients are required to include the *Content-Length* HTTP header with the correct content length in every request that can have a body (e.g. *POST*, *PUT* or *PATCH*) request. ArangoDB will not process requests without a *Content-Length* header - thus chunked transfer encoding for POST-documents is not supported.

HTTP Keep-Alive

ArangoDB supports HTTP keep-alive. If the client does not send a *Connection* header in its request, and the client uses HTTP version 1.1, ArangoDB will assume the client wants to keep alive the connection. If clients do not wish to use the keep-alive feature, they should explicitly indicate that by sending a *Connection: Close* HTTP header in the request.

ArangoDB will close connections automatically for clients that send requests using HTTP 1.0, except if they send an *Connection: Keep-Alive* header.

The default Keep-Alive timeout can be specified at server start using the *--http.keep-alive-timeout* parameter.

Establishing TCP connections is expensive, since it takes several ping pongs between the communication parties. Therefore you can use connection keepalive to send several HTTP request over one TCP-connection; Each request is treated independently by definition. You can use this feature to build up a so called *connection pool* with several established connections in your client application, and dynamically re-use one of those then idle connections for subsequent requests.

Blocking vs. Non-blocking HTTP Requests

ArangoDB supports both blocking and non-blocking HTTP requests.

ArangoDB is a multi-threaded server, allowing the processing of multiple client requests at the same time. Request/response handling and the actual work are performed on the server in parallel by multiple worker threads.

Still, clients need to wait for their requests to be processed by the server, and thus keep one connection of a pool occupied. By default, the server will fully process an incoming request and then return the result to the client when the operation is finished. The client must wait for the server's HTTP response before it can send additional requests over the same connection. For clients that are single-threaded and/or are blocking on I/O themselves, waiting idle for the server response may be non-optimal.

To reduce blocking on the client side, ArangoDB offers a generic mechanism for non-blocking, asynchronous execution: clients can add the HTTP header *x-arango-async: true* to any of their requests, marking them as to be executed asynchronously on the server. ArangoDB will put such requests into an in-memory task queue and return an *HTTP 202* (accepted) response to the client instantly and thus finish this HTTP-request. The server will execute the tasks from the queue asynchronously as fast as possible, while clients can continue to do other work. If the server queue is full (i.e. contains as many tasks as specified by the option "--scheduler.maximal-queue-size"), then the request will be rejected instantly with an *HTTP 500* (internal server error) response.

Asynchronous execution decouples the request/response handling from the actual work to be performed, allowing fast server responses and greatly reducing wait time for clients. Overall this allows for much higher throughput than if clients would always wait for the server's response.

Keep in mind that the asynchronous execution is just "fire and forget". Clients will get any of their asynchronous requests answered with a generic HTTP 202 response. At the time the server sends this response, it does not know whether the requested operation can be carried out successfully (the actual operation execution will happen at some later point). Clients therefore cannot make a decision based on the server response and must rely on their requests being valid and processable by the server.

Additionally, the server's asynchronous task queue is an in-memory data structure, meaning not-yet processed tasks from the queue might be lost in case of a crash. Clients should therefore not use the asynchronous feature when they have strict durability requirements or if they rely on the immediate result of the request they send.

For details on the subsequent processing read on under Async Result handling.

Authentication

Client authentication can be achieved by using the *Authorization* HTTP header in client requests. ArangoDB supports authentication via HTTP Basic or JWT.

Authentication is turned on by default, but can be turned off. To enforce authentication for incoming requested, the server must be started with the option --server.authentication. Please note that requests using the HTTP OPTIONS method will be answered by ArangoDB in any case, even if no authentication data is sent by the client or if the authentication data is wrong. This is required for handling CORS preflight requests (see Cross Origin Resource Sharing requests). The response to an HTTP OPTIONS request will be generic and not expose any private data.

Please note that when authentication is turned on in ArangoDB, it will by default affect all incoming requests.

There is an additional option --server.authentication-system-only to restrict authentication to requests to the ArangoDB internal APIs and the admin interface. This option can be used to expose a public API built with ArangoDB to the outside world without the need for HTTP authentication, but to still protect the usage of the ArangoDB API (i.e. */_api/*) and the admin interface (i.e. */_admin/*) with HTTP authentication.

If the server is started with the *--server.authentication-system-only* parameter set to *false* (which is the default), all incoming requests need HTTP authentication if the server is configured to require HTTP authentication. Setting the option to *false* will make the server require authentication only for requests to the internal functionality at */_api/* or */_admin* and will allow unauthenticated requests to all other URLs.

Whenever authentication is required and the client has not yet authenticated, ArangoDB will return *HTTP 401* (Unauthorized). It will also send the *WWW-Authenticate* response header, indicating that the client should prompt the user for username and password if supported. If the client is a browser, then sending back this header will normally trigger the display of the browser-side HTTP authentication dialog. As showing the browser HTTP authentication dialog is undesired in AJAX requests, ArangoDB can be told to not send the *WWW-Authenticate* header back to the client. Whenever a client sends the *X-Omit-WWW-Authenticate* HTTP header (with an arbitrary value) to ArangoDB, ArangoDB will only send status code 401, but no *WWW-Authenticate* header. This allows clients to implement credentials handling and bypassing the browser's built-in dialog.

Authentication via JWT

To authenticate via JWT you must first obtain a JWT. To do so send a POST request to

/_open/auth

containing *username* and *password* JSON-encoded like so:

{"username":"root","password":"rootPassword"}

Upon success the endpoint will return a 200 OK and an answer containing the JWT in a JSON- encoded object like so:

```
{"jwt":"eyJhbGciOiJIUzI1NiI..x6EfI","must_change_password":false}
```

This JWT should then be used within the Authorization HTTP header in subsequent requests:

```
Authorization: bearer eyJhbGciOiJIUzI1NiI..x6EfI
```

Please note that the JWT will expire after 1 month and needs to be updated.

ArangoDB uses a standard JWT authentication. The secret may either be set using `--server.jwt-secret` or will be randomly generated upon server startup.

For more information on JWT please consult RFC7519 and https://jwt.io

Error Handling

The following should be noted about how ArangoDB handles client errors in its HTTP layer:

- client requests using an HTTP version signature different than *HTTP/1.0* or *HTTP/1.1* will get an *HTTP 505* (HTTP version not supported) error in return.
- ArangoDB will reject client requests with a negative value in the *Content-Length* request header with *HTTP 411* (Length Required).
- Arangodb doesn't support POST with *transfer-encoding: chunked* which forbids the *Content-Length* header above.
- the maximum URL length accepted by ArangoDB is 16K. Incoming requests with longer URLs will be rejected with an *HTTP 414* (Request-URI too long) error.
- if the client sends a *Content-Length* header with a value bigger than 0 for an HTTP GET, HEAD, or DELETE request, ArangoDB will process the request, but will write a warning to its log file.
- when the client sends a *Content-Length* header that has a value that is lower than the actual size of the body sent, ArangoDB will respond with *HTTP 400* (Bad Request).
- if clients send a *Content-Length* value bigger than the actual size of the body of the request, ArangoDB will wait for about 90 seconds for the client to complete its request. If the client does not send the remaining body data within this time, ArangoDB will close the connection. Clients should avoid sending such malformed requests as this will block one tcp connection, and may lead to a temporary filedescriptor leak.
- when clients send a body or a *Content-Length* value bigger than the maximum allowed value (512 MB), ArangoDB will respond with *HTTP 413* (Request Entity Too Large).
- if the overall length of the HTTP headers a client sends for one request exceeds the maximum allowed size (1 MB), the server will fail with *HTTP 431* (Request Header Fields Too Large).
- if clients request an HTTP method that is not supported by the server, ArangoDB will return with *HTTP 405* (Method Not Allowed). ArangoDB offers general support for the following HTTP methods:
  - GET
  - POST
  - PUT
  - DELETE
  - HEAD
  - PATCH
  - OPTIONS

  Please note that not all server actions allow using all of these HTTP methods. You should look up up the supported methods for each method you intend to use in the manual.

  Requests using any other HTTP method (such as for example CONNECT, TRACE etc.) will be rejected by ArangoDB as mentioned before.

Cross Origin Resource Sharing (CORS) requests

ArangoDB will automatically handle CORS requests as follows:

- when the client sends an *Origin* HTTP header, ArangoDB will return a header *access-control-allow-origin* containing the value the client sent in the *Origin* header.
- for non-trivial CORS requests, clients may issue a preflight request via an additional HTTP OPTIONS request. ArangoDB will automatically answer such preflight HTTP OPTIONS requests with an HTTP 200 response with an empty body. ArangoDB will return the following headers in the response:
  - *access-control-allow-origin*: will contain the value that the client provided in the *Origin* header of the request

  - *access-control-allow-methods*: will contain an array of all HTTP methods generally supported by ArangoDB. This array does not depend on the URL the client requested and is the same for all CORS requests.
  - *access-control-allow-headers*: will contain exactly the value that the client has provided in the *Access-Control-Request-Header* header of the request. This header will only be returned if the client has specified the header in the request. ArangoDB will send back the original value without further validation.
  - *access-control-max-age*: will return a cache lifetime for the preflight response as determined by ArangoDB.
- any *access-control-allow-credentials* header sent by the client is ignored by ArangoDB if its value is not *true*. If a client sends a header value of *true*, ArangoDB will return the header *access-control-allow-credentials: true* too, but only if the value of the sent `Origin` header matches a trusted origin in the `--http.trusted-origin` startup option. To make ArangoDB trust a certain origin, specify the origin at server start like this: `--http.trusted-origin "http://localhost:8529"` To specify multiple trusted origins, the option can be specified multiple times. To trust any origin, the special value `*` can be specified as a trusted origin: `--http.trusted-origin "*"`

Note that CORS preflight requests will probably not send any authentication data with them. One of the purposes of the preflight request is to check whether the server accepts authentication or not.

A consequence of this is that ArangoDB will allow requests using the HTTP OPTIONS method without credentials, even when the server is run with authentication enabled.

The response to the HTTP OPTIONS request will however be a generic response that will not expose any private data and thus can be considered "safe" even without credentials.

HTTP method overriding

ArangoDB provides a startup option *--http.allow-method-override*. This option can be set to allow overriding the HTTP request method (e.g. GET, POST, PUT, DELETE, PATCH) of a request using one of the following custom HTTP headers:

- *x-http-method-override*
- *x-http-method*
- *x-method-override*

This allows using HTTP clients that do not support all "common" HTTP methods such as PUT, PATCH and DELETE. It also allows bypassing proxies and tools that would otherwise just let certain types of requests (e.g. GET and POST) pass through.

Enabling this option may impose a security risk, so it should only be used in very controlled environments. Thus the default value for this option is *false* (no method overriding allowed). You need to enable it explicitly if you want to use this feature.

HTTP Interface

Following you have ArangoDB's HTTP Interface for Documents, Databases, Edges and more.

There are also some examples provided for every API action.

You may also use the interactive Swagger documentation in the ArangoDB webinterface to explore the API calls below.

HTTP Interface for Databases

Address of a Database

Any operation triggered via ArangoDB's HTTP REST API is executed in the context of exactly one database. To explicitly specify the database in a request, the request URI must contain the database name in front of the actual path:

```
http://localhost:8529/_db/mydb/...
```

where ... is the actual path to the accessed resource. In the example, the resource will be accessed in the context of the database *mydb*. Actual URLs in the context of *mydb* could look like this:

```
http://localhost:8529/_db/mydb/_api/version
http://localhost:8529/_db/mydb/_api/document/test/12345
http://localhost:8529/_db/mydb/myapp/get
```

Database-to-Endpoint Mapping

If a database name is present in the URI as above, ArangoDB will consult the database-to-endpoint mapping for the current endpoint, and validate if access to the database is allowed on the endpoint. If the endpoint is not restricted to an array of databases, ArangoDB will continue with the regular authentication procedure. If the endpoint is restricted to an array of specified databases, ArangoDB will check if the requested database is in the array. If not, the request will be turned down instantly. If yes, then ArangoDB will continue with the regular authentication procedure.

If the request URI was *http:// localhost:8529/_db/mydb/...*, then the request to *mydb* will be allowed (or disallowed) in the following situations:

```
Endpoint-to-database mapping           Access to *mydb* allowed?
----------------------------           -------------------------
[ ]                                    yes
[ "_system" ]                          no
[ "_system", "mydb" ]                  yes
[ "mydb" ]                             yes
[ "mydb", "_system" ]                  yes
[ "test1", "test2" ]                   no
```

In case no database name is specified in the request URI, ArangoDB will derive the database name from the endpoint-to-database mapping of the endpoint the connection was coming in on.

If the endpoint is not restricted to an array of databases, ArangoDB will assume the *_system* database. If the endpoint is restricted to one or multiple databases, ArangoDB will assume the first name from the array.

Following is an overview of which database name will be assumed for different endpoint-to-database mappings in case no database name is specified in the URI:

```
Endpoint-to-database mapping           Database
----------------------------           --------
[ ]                                    _system
[ "_system" ]                          _system
[ "_system", "mydb" ]                  _system
[ "mydb" ]                             mydb
[ "mydb", "_system" ]                  mydb
```

Database Management

This is an introduction to ArangoDB's HTTP interface for managing databases.

The HTTP interface for databases provides operations to create and drop individual databases. These are mapped to the standard HTTP methods *POST* and *DELETE*. There is also the *GET* method to retrieve an array of existing databases.

Please note that all database management operations can only be accessed via the default database (*_system*) and none of the other databases.

Managing Databases using HTTP

Information of the database

retrieves information about the current database

`GET /_api/database/current`

Retrieves information about the current database

The response is a JSON object with the following attributes:

- *name*: the name of the current database
- *id*: the id of the current database
- *path*: the filesystem path of the current database
- *isSystem*: whether or not the current database is the *_system* database

Example:

```
shell> curl --dump - http://localhost:8529/_api/database/current

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "result" : {
    "name" : "_system",
    "id" : "1",
    "path" : "/tmp/arangosh_GgAPmL/tmp-15793-280372105/data/databases/database-1",
    "isSystem" : true
  },
  "error" : false,
  "code" : 200
}
```

Return Codes

- *200*: is returned if the information was retrieved successfully.
- *400*: is returned if the request is invalid.
- *404*: is returned if the database could not be found.

Examples

```
shell> curl --dump - http://localhost:8529/_api/database/current

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

List of accessible databases

retrieves a list of all databases the current user can access

`GET /_api/database/user`

Retrieves the list of all databases the current user can access without specifying a different username or password.

Example:

```
shell> curl --dump - http://localhost:8529/_api/database/user

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "result" : [
    "_system"
  ],
  "error" : false,
  "code" : 200
}
```

Return Codes

- *200*: is returned if the list of database was compiled successfully.
- *400*: is returned if the request is invalid.

Examples

```
shell> curl --dump - http://localhost:8529/_api/database/user

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

List of databases

retrieves a list of all existing databases

`GET /_api/database`

Retrieves the list of all existing databases

Note: retrieving the list of databases is only possible from within the *_system* database.

Note: You should use the *GET user API* to fetch the list of the available databases now.

Example:

```
shell> curl --dump - http://localhost:8529/_api/database

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "result" : [
    "_system"
  ],
  "error" : false,
  "code" : 200
}
```

Return Codes

- *200*: is returned if the list of database was compiled successfully.
- *400*: is returned if the request is invalid.
- *403*: is returned if the request was not executed in the *_system* database.

Examples

```
shell> curl --dump - http://localhost:8529/_api/database

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Create database

creates a new database

`POST /_api/database`

A json post document with these Properties is required:

- **name**: Has to contain a valid database name.
- **users**: Has to be an array of user objects to initially create for the new database. User information will not be changed for users that already exist. If *users* is not specified or does not contain any users, a default user *root* will be created with an empty string password. This ensures that the new database will be accessible after it is created. Each user object can contain the following attributes:
  - **username**: Loginname of the user to be created
  - **passwd**: The user password as a string. If not specified, it will default to an empty string.
  - **active**: A flag indicating whether the user account should be activated or not. The default value is *true*. If set to *false*, the user won't be able to log into the database.
  - **extra**: A JSON object with extra user information. The data contained in *extra* will be stored for the user but not be interpreted further by ArangoDB.

Creates a new database

The response is a JSON object with the attribute *result* set to *true*.

Note: creating a new database is only possible from within the *_system* database.

Example:

Creating a database named *example*.

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/database <<EOF
{
  "name" : "example"
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8

{
  "result" : true,
  "error" : false,
  "code" : 201
}
```

Example:

Creating a database named *mydb* with two users.

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/database <<EOF
{
  "name" : "mydb",
  "users" : [
    {
      "username" : "admin",
      "passwd" : "secret",
      "active" : true
    },
    {
      "username" : "tester",
      "passwd" : "test001",
      "active" : false
    }
  ]
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8

{
  "result" : true,
  "error" : false,
  "code" : 201
}
```


```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/database <<EOF
{
  "name" : "example"
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8
```


show response body

Creating a database named *mydb* with two users.

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/database <<EOF
{
  "name" : "mydb",
  "users" : [
    {
      "username" : "admin",
      "passwd" : "secret",
      "active" : true
    },
    {
      "username" : "tester",
      "passwd" : "test001",
      "active" : false
    }
  ]
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8
```

show response body

Drop database

drop an existing database

`DELETE /_api/database/{database-name}`

Path Parameters

- *database-name* (required): The name of the database

Drops the database along with all data stored in it.

Note: dropping a database is only possible from within the *_system* database. The *_system* database itself cannot be dropped.

Example:

```
shell> curl -X DELETE --dump - http://localhost:8529/_api/database/example

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "result" : true,
  "error" : false,
  "code" : 200
}
```

Return Codes

- *200*: is returned if the database was dropped successfully.
- *400*: is returned if the request is malformed.
- *403*: is returned if the request was not executed in the *_system* database.
- *404*: is returned if the database could not be found.

Examples

```
shell> curl -X DELETE --dump - http://localhost:8529/_api/database/example

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Notes on Databases

Please keep in mind that each database contains its own system collections, which need to set up when a database is created. This will make the creation of a database take a while. Replication is configured on a per-database level, meaning that any replication logging or applying for the a new database must be configured explicitly after a new database has been created. Foxx applications are also available only in the context of the database they have been installed in. A new database will only provide access to the system applications shipped with ArangoDB (that is the web interface at the moment) and no other Foxx applications until they are explicitly installed for the particular database.

Database

ArangoDB can handle multiple databases in the same server instance. Databases can be used to logically group and separate data. An ArangoDB database consists of collections and dedicated database-specific worker processes. A database contains its own collections (which cannot be accessed from other databases), Foxx applications and replication loggers and appliers. Each ArangoDB database contains its own system collections (e.g. _users, _graphs, ...).

There will always be at least one database in ArangoDB. This is the default database named _system. This database cannot be dropped and provides special operations for creating, dropping and enumerating databases. Users can create additional databases and give them unique names to access them later. Database management operations cannot be initiated from out of user-defined databases.

When ArangoDB is accessed via its HTTP REST API, the database name is read from the first part of the request URI path (e.g. /_db/_system/...). If the request URI does not contain a database name, the database name is automatically determined by the algorithm described in Database-to-Endpoint Mapping .

Database Name

A single ArangoDB instance can handle multiple databases in parallel. When multiple databases are used, each database must be given an unique name. This name is used to uniquely identify a database. The default database in ArangoDB is named *system. The database name is a string consisting of only letters, digits and the* (underscore) and - (dash) characters. User-defined database names must always start with a letter. Database names are case-sensitive.

Database Organization

A single ArangoDB instance can handle multiple databases in parallel. By default, there will be at least one database which is named _system. Databases are physically stored in separate sub-directories underneath the database directory, which itself resides in the instance's data directory.

Each database has its own sub-directory, named database-. The database directory contains sub-directories for the collections of the database, and a file named parameter.json. This file contains the database id and name.

In an example ArangoDB instance which has two databases, the filesystem layout could look like this:

```
data/                     # the instance's data directory
  databases/              # sub-directory containing all databases' data
    database-<id>/        # sub-directory for a single database
      parameter.json      # file containing database id and name
      collection-<id>/    # directory containing data about a collection
    database-<id>/        # sub-directory for another database
      parameter.json      # file containing database id and name
      collection-<id>/    # directory containing data about a collection
      collection-<id>/    # directory containing data about a collection
```

Foxx applications are also organized in database-specific directories inside the application path. The filesystem layout could look like this:

```
apps/                   # the instance's application directory
  system/               # system applications (can be ignored)
  databases/            # sub-directory containing database-specific applications
    <database-name>/    # sub-directory for a single database
      <app-name>        # sub-directory for a single application
      <app-name>        # sub-directory for a single application
    <database-name>/    # sub-directory for another database
      <app-name>        # sub-directory for a single application
`
```

HTTP Interface for Collections

Collections

This is an introduction to ArangoDB's HTTP interface for collections.

Collection

A collection consists of documents. It is uniquely identified by its collection identifier. It also has a unique name that clients should use to identify and access it. Collections can be renamed. This will change the collection name, but not the collection identifier. Collections have a type that is specified by the user when the collection is created. There are currently two types: document and edge. The default type is document.

Collection Identifier

A collection identifier lets you refer to a collection in a database. It is a string value and is unique within the database. Up to including ArangoDB 1.1, the collection identifier has been a client's primary means to access collections. Starting with ArangoDB 1.2, clients should instead use a collection's unique name to access a collection instead of its identifier. ArangoDB currently uses 64bit unsigned integer values to maintain collection ids internally. When returning collection ids to clients, ArangoDB will put them into a string to ensure the collection id is not clipped by clients that do not support big integers. Clients should treat the collection ids returned by ArangoDB as opaque strings when they store or use it locally.

Note: collection ids have been returned as integers up to including ArangoDB 1.1

Collection Name

A collection name identifies a collection in a database. It is a string and is unique within the database. Unlike the collection identifier it is supplied by the creator of the collection. The collection name must consist of letters, digits, and the _ (underscore) and - (dash) characters only. Please refer to Naming Conventions in ArangoDB for more information on valid collection names.

Key Generator

ArangoDB allows using key generators for each collection. Key generators have the purpose of auto-generating values for the _key attribute of a document if none was specified by the user. By default, ArangoDB will use the traditional key generator. The traditional key generator will auto-generate key values that are strings with ever-increasing numbers. The increment values it uses are non-deterministic.

Contrary, the auto increment key generator will auto-generate deterministic key values. Both the start value and the increment value can be defined when the collection is created. The default start value is 0 and the default increment is 1, meaning the key values it will create by default are:

1, 2, 3, 4, 5, ...

When creating a collection with the auto increment key generator and an increment of 5, the generated keys would be:

1, 6, 11, 16, 21, ...

The basic operations (create, read, update, delete) for documents are mapped to the standard HTTP methods (*POST*, *GET*, *PUT*, *DELETE*).

Address of a Collection

All collections in ArangoDB have an unique identifier and a unique name. ArangoDB internally uses the collection's unique identifier to look up collections. This identifier however is managed by ArangoDB and the user has no control over it. In order to allow users use their own names, each collection also has a unique name, which is specified by the user. To access a collection from the user perspective, the collection name should be used, i.e.:

```
http://server:port/_api/collection/collection-name
```

For example: Assume that the collection identifier is *7254820* and the collection name is *demo*, then the URL of that collection is:

```
http://localhost:8529/_api/collection/demo
```

Creating and Deleting Collections

Create collection

creates a collection

`POST /_api/collection`

A json post document with these Properties is required:

- **journalSize**: The maximal size of a journal or datafile in bytes. The value must be at least `1048576` (1 MiB). (The default is a configuration parameter)
- **replicationFactor**: (The default is *1*): in a cluster, this attribute determines how many copies of each shard are kept on different DBServers. The value 1 means that only one copy (no synchronous replication) is kept. A value of k means that k-1 replicas are kept. Any two copies reside on different DBServers. Replication between them is synchronous, that is, every write operation to the "leader" copy will be replicated to all "follower" replicas, before the write operation is reported successful. If a server fails, this is detected automatically and one of the servers holding copies take over, usually without an error being reported.
- **keyOptions**:
  - **allowUserKeys**: if set to *true*, then it is allowed to supply own key values in the *_key* attribute of a document. If set to *false*, then the key generator will solely be responsible for generating keys and supplying own key values in the *_key* attribute of documents is considered an error.
  - **type**: specifies the type of the key generator. The currently available generators are *traditional* and *autoincrement*.
  - **increment**: increment value for *autoincrement* key generator. Not used for other key generator types.
  - **offset**: Initial offset value for *autoincrement* key generator. Not used for other key generator types.
- **name**: The name of the collection.
- **waitForSync**: If *true* then the data is synchronized to disk before returning from a document create, update, replace or removal operation. (default: false)
- **doCompact**: whether or not the collection will be compacted (default is *true*)
- **isVolatile**: If *true* then the collection data is kept in-memory only and not made persistent. Unloading the collection will cause the collection data to be discarded. Stopping or re-starting the server will also cause full loss of data in the collection. Setting this option will make the resulting collection be slightly faster than regular collections because ArangoDB does not enforce any synchronization to disk and does not calculate any CRC checksums for datafiles (as there are no datafiles). This option should therefore be used for cache-type collections only, and not for data that cannot be re-created otherwise. (The default is *false*)
- **shardKeys**: (The default is *["_key"]*): in a cluster, this attribute determines which document attributes are used to determine the target shard for documents. Documents are sent to shards based on the values of their shard key attributes. The values of all shard key attributes in a document are hashed, and the hash value is used to determine the target shard. **Note**: Values of shard key attributes cannot be changed once set. This option is meaningless in a single server setup.
- **numberOfShards**: (The default is *1*): in a cluster, this value determines the number of shards to create for the collection. In a single server setup, this option is meaningless.
- **isSystem**: If *true*, create a system collection. In this case *collection-name* should start with an underscore. End users should normally create non-system collections only. API implementors may be required to create system collections in very special occasions, but normally a regular collection will do. (The default is *false*)
- **type**: (The default is *2*): the type of the collection to create. The following values for *type* are valid:
  - *2*: document collection
  - *3*: edges collection
- **indexBuckets**: The: number of buckets into which indexes using a hash table are split. The default is 16 and this number has to be a power of 2 and less than or equal to 1024. For very large collections one should increase this to avoid long pauses when the hash table has to be initially built or resized, since buckets are resized individually and can be initially built in parallel. For example, 64 might be a sensible value for a collection with 100 000 000 documents. Currently, only the edge index respects this value, but other index types might follow in future ArangoDB versions. Changes (see below) are applied when the collection is loaded the next time.

Creates a new collection with a given name. The request must contain an object with the following attributes.

Example:

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/collection <<EOF
{
  "name" : "testCollectionBasics"
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
location: /_db/_system/_api/collection/testCollectionBasics

{
  "id" : "9107",
  "name" : "testCollectionBasics",
  "waitForSync" : false,
  "isVolatile" : false,
  "isSystem" : false,
  "status" : 3,
  "type" : 2,
  "error" : false,
  "code" : 200
}
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/collection <<EOF
{
  "name" : "testCollectionEdges",
  "type" : 3
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
location: /_db/_system/_api/collection/testCollectionEdges

{
  "id" : "9109",
  "name" : "testCollectionEdges",
  "waitForSync" : false,
  "isVolatile" : false,
  "isSystem" : false,
  "status" : 3,
  "type" : 3,
  "error" : false,
  "code" : 200
}
```

Example:

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/collection <<EOF
{
  "name" : "testCollectionUsers",
  "keyOptions" : {
    "type" : "autoincrement",
    "increment" : 5,
    "allowUserKeys" : true
  }
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
location: /_db/_system/_api/collection/testCollectionUsers

{
  "id" : "9113",
  "name" : "testCollectionUsers",
  "waitForSync" : false,
  "isVolatile" : false,
  "isSystem" : false,
  "status" : 3,
  "type" : 2,
  "error" : false,
  "code" : 200
}
```

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/collection
<<EOF
{
  "name" : "testCollectionBasics"
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
location: /_db/_system/_api/collection/testCollectionBasics
```

show response body

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/collection
<<EOF
{
  "name" : "testCollectionUsers",
  "keyOptions" : {
    "type" : "autoincrement",
    "increment" : 5,
    "allowUserKeys" : true
  }
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
location: /_db/_system/_api/collection/testCollectionUsers
```

show response body

Drops a collection

drops a collection

`DELETE /_api/collection/{collection-name}`

Path Parameters

- *collection-name* (required): The name of the collection to drop.

Query Parameters

- *isSystem* (optional): Whether or not the collection to drop is a system collection. This parameter must be set to *true* in order to drop a system collection.

Drops the collection identified by *collection-name*.

If the collection was successfully dropped, an object is returned with the following attributes:

- *error*: *false*
- *id*: The identifier of the dropped collection.

Example:

Using an identifier:

```
shell> curl -X DELETE --dump - http://localhost:8529/_api/collection/9116

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "id" : "9116",
  "error" : false,
  "code" : 200
}
```

Example:

Using a name:

```
shell> curl -X DELETE --dump - http://localhost:8529/_api/collection/products1

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "id" : "9119",
  "error" : false,
  "code" : 200
}
```

Example:

Dropping a system collection

```
shell> curl -X DELETE --dump - http://localhost:8529/_api/collection/_example?isSystem=true

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "id" : "9122",
  "error" : false,
  "code" : 200
}
```

Return Codes

- *400*: If the *collection-name* is missing, then a *HTTP 400* is returned.
- *404*: If the *collection-name* is unknown, then a *HTTP 404* is returned.

Examples

Using an identifier:

```
shell> curl -X DELETE --dump - http://localhost:8529/_api/collection/9116

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body
Using a name:

```
shell> curl -X DELETE --dump - http://localhost:8529/_api/collection/products1

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body
Dropping a system collection

```
shell> curl -X DELETE --dump - http://localhost:8529/_api/collection/_example?
isSystem=true

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Truncate collection

truncates a collection

`PUT /_api/collection/{collection-name}/truncate`

Path Parameters

- *collection-name* (required): The name of the collection.

Removes all documents from the collection, but leaves the indexes intact.

Example:

```
shell> curl -X PUT --dump - http://localhost:8529/_api/collection/products/truncate

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "id" : "9500",
  "name" : "products",
  "isSystem" : false,
  "status" : 3,
  "type" : 2,
  "error" : false,
  "code" : 200
}
```

Return Codes

- *400*: If the *collection-name* is missing, then a *HTTP 400* is returned.
- *404*: If the *collection-name* is unknown, then a *HTTP 404* is returned.

Examples

```
shell> curl -X PUT --dump - http://localhost:8529/_api/collection/products/truncate

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Getting Information about a Collection

Return information about a collection

returns a collection

`GET /_api/collection/{collection-name}`

Path Parameters

- *collection-name* (required): The name of the collection.

The result is an object describing the collection with the following attributes:

- *id*: The identifier of the collection.
- *name*: The name of the collection.
- *status*: The status of the collection as number.
  - 1: new born collection
  - 2: unloaded
  - 3: loaded
  - 4: in the process of being unloaded
  - 5: deleted
  - 6: loading

Every other status indicates a corrupted collection.

- *type*: The type of the collection as number.
  - 2: document collection (normal case)
  - 3: edges collection
- *isSystem*: If *true* then the collection is a system collection.

Return Codes

- *404*: If the *collection-name* is unknown, then a *HTTP 404* is returned.

Read properties of a collection

reads the properties of the specified collection

`GET /_api/collection/{collection-name}/properties`

Path Parameters

- *collection-name* (required): The name of the collection.

In addition to the above, the result will always contain the *waitForSync*, *doCompact*, *journalSize*, and *isVolatile* attributes. This is achieved by forcing a load of the underlying collection.

- *waitForSync*: If *true* then creating, changing or removing documents will wait until the data has been synchronized to disk.
- *doCompact*: Whether or not the collection will be compacted.
- *journalSize*: The maximal size setting for journals / datafiles in bytes.
- *keyOptions*: JSON object which contains key generation options:
  - *type*: specifies the type of the key generator. The currently available generators are *traditional* and *autoincrement*.
  - *allowUserKeys*: if set to *true*, then it is allowed to supply own key values in the *_key* attribute of a document. If set to *false*, then the key generator is solely responsible for generating keys and supplying own key values in the *_key* attribute of documents is considered an error.
- *isVolatile*: If *true* then the collection data will be kept in memory only and ArangoDB will not write or sync the data to disk.

In a cluster setup, the result will also contain the following attributes:

- *numberOfShards*: the number of shards of the collection.
- *shardKeys*: contains the names of document attributes that are used to determine the target shard for documents.
- *replicationFactor*: contains how many copies of each shard are kept on different DBServers.

Example:

Using an identifier:

```
shell> curl --dump - http://localhost:8529/_api/collection/9472/properties

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
location: /_db/_system/_api/collection/products/properties

{
  "id" : "9472",
  "name" : "products",
  "isSystem" : false,
  "doCompact" : true,
  "isVolatile" : false,
  "journalSize" : 33554432,
  "keyOptions" : {
    "type" : "traditional",
    "allowUserKeys" : true
  },
  "waitForSync" : true,
  "indexBuckets" : 8,
  "status" : 3,
  "type" : 2,
  "error" : false,
  "code" : 200
}
```

Example:

Using a name:

```
shell> curl --dump - http://localhost:8529/_api/collection/products/properties

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
location: /_db/_system/_api/collection/products/properties

{
  "id" : "9476",
  "name" : "products",
  "isSystem" : false,
  "doCompact" : true,
  "isVolatile" : false,
  "journalSize" : 33554432,
  "keyOptions" : {
    "type" : "traditional",
    "allowUserKeys" : true
  },
  "waitForSync" : true,
  "indexBuckets" : 8,
  "status" : 3,
  "type" : 2,
  "error" : false,
  "code" : 200
}
```

Return Codes

- *400*: If the *collection-name* is missing, then a *HTTP 400* is returned.
- *404*: If the *collection-name* is unknown, then a *HTTP 404* is returned.

Examples

Using an identifier:

```
shell> curl --dump - http://localhost:8529/_api/collection/9472/properties

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
location: /_db/_system/_api/collection/products/properties
```

show response body
Using a name:

```
shell> curl --dump - http://localhost:8529/_api/collection/products/properties

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
location: /_db/_system/_api/collection/products/properties
```

show response body

Return number of documents in a collection

Counts the documents in a collection

`GET /_api/collection/{collection-name}/count`

Path Parameters

- *collection-name* (required): The name of the collection.

In addition to the above, the result also contains the number of documents. **Note** that this will always load the collection into memory.

- *count*: The number of documents inside the collection.

Example:

Requesting the number of documents:

```
shell> curl --dump - http://localhost:8529/_api/collection/products/count

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
location: /_db/_system/_api/collection/products/count

{
  "id" : "9142",
  "name" : "products",
  "isSystem" : false,
  "doCompact" : true,
  "isVolatile" : false,
  "journalSize" : 33554432,
  "keyOptions" : {
    "type" : "traditional",
    "allowUserKeys" : true
  },
  "waitForSync" : true,
  "indexBuckets" : 8,
  "count" : 100,
  "status" : 3,
  "type" : 2,
  "error" : false,
  "code" : 200
}
```

Return Codes

- *400*: If the *collection-name* is missing, then a *HTTP 400* is returned.
- *404*: If the *collection-name* is unknown, then a *HTTP 404* is returned.

Examples

Requesting the number of documents:

```
shell> curl --dump - http://localhost:8529/_api/collection/products/count

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
location: /_db/_system/_api/collection/products/count
```

show response body

Return statistics for a collection

Fetch the statistics of a collection

`GET /_api/collection/{collection-name}/figures`

Path Parameters

- *collection-name* (required): The name of the collection.

In addition to the above, the result also contains the number of documents and additional statistical information about the collection. **Note** : This will always load the collection into memory.

Note: collection data that are stored in the write-ahead log only are not reported in the results. When the write-ahead log is collected, documents might be added to journals and datafiles of the collection, which may modify the figures of the collection.

Additionally, the filesizes of collection and index parameter JSON files are not reported. These files should normally have a size of a few bytes each. Please also note that the *fileSize* values are reported in bytes and reflect the logical file sizes. Some filesystems may use optimisations (e.g. sparse files) so that the actual physical file size is somewhat different. Directories and sub-directories may also require space in the file system, but this space is not reported in the *fileSize* results.

That means that the figures reported do not reflect the actual disk usage of the collection with 100% accuracy. The actual disk usage of a collection is normally slightly higher than the sum of the reported *fileSize* values. Still the sum of the *fileSize* values can still be used as a lower bound approximation of the disk usage.

A json document with these Properties is returned:

HTTP 200

Returns information about the collection:

- **count**: The number of documents currently present in the collection.
- **journalSize**: The maximal size of a journal or datafile in bytes.
- **figures**:
  - **datafiles**:
    - **count**: The number of datafiles.
    - **fileSize**: The total filesize of datafiles (in bytes).
  - **uncollectedLogfileEntries**: The number of markers in the write-ahead log for this collection that have not been transferred to journals or datafiles.
  - **documentReferences**: The number of references to documents in datafiles that JavaScript code currently holds. This information can be used for debugging compaction and unload issues.
  - **compactionStatus**:
    - **message**: The action that was performed when the compaction was last run for the collection. This information can be used for debugging compaction issues.
    - **time**: The point in time the compaction for the collection was last executed. This information can be used for debugging compaction issues.
  - **compactors**:
    - **count**: The number of compactor files.
    - **fileSize**: The total filesize of all compactor files (in bytes).
  - **dead**:
    - **count**: The number of dead documents. This includes document versions that have been deleted or replaced by a newer version. Documents deleted or replaced that are contained the write-ahead log only are not reported in this figure.
    - **deletion**: The total number of deletion markers. Deletion markers only contained in the write-ahead log are not reporting in this figure.
    - **size**: The total size in bytes used by all dead documents.
  - **indexes**:
    - **count**: The total number of indexes defined for the collection, including the pre-defined indexes (e.g. primary index).
    - **size**: The total memory allocated for indexes in bytes.
  - **readcache**:
    - **count**: The number of revisions of this collection stored in the document revisions cache.
    - **size**: The memory used for storing the revisions of this collection in the document revisions cache (in bytes). This figure does not include the document data but only mappings from document revision ids to cache entry locations.
  - **waitingFor**: An optional string value that contains information about which object type is at the head of the collection's cleanup queue. This information can be used for debugging compaction and unload issues.
  - **alive**:
    - **count**: The number of currently active documents in all datafiles and journals of the collection. Documents that are contained in the write-ahead log only are not reported in this figure.
    - **size**: The total size in bytes used by all active documents of the collection. Documents that are contained in the write-ahead log only are not reported in this figure.
  - **lastTick**: The tick of the last marker that was stored in a journal of the collection. This might be 0 if the collection does not yet have a journal.
  - **journals**:
    - **count**: The number of journal files.
    - **fileSize**: The total filesize of all journal files (in bytes).
  - **revisions**:
    - **count**: The number of revisions of this collection managed by the storage engine.
    - **size**: The memory used for storing the revisions of this collection in the storage engine (in bytes). This figure does not

include the document data but only mappings from document revision ids to storage engine datafile positions.

Example:

Using an identifier and requesting the figures of the collection:

```
shell> curl --dump - http://localhost:8529/_api/collection/products/figures

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
location: /_db/_system/_api/collection/products/figures

{
  "id" : "9448",
  "name" : "products",
  "isSystem" : false,
  "doCompact" : true,
  "isVolatile" : false,
  "journalSize" : 33554432,
  "keyOptions" : {
    "type" : "traditional",
    "allowUserKeys" : true
  },
  "waitForSync" : false,
  "indexBuckets" : 8,
  "count" : 1,
  "figures" : {
    "indexes" : {
      "count" : 1,
      "size" : 32128
    },
    "lastTick" : 9453,
    "uncollectedLogfileEntries" : 0,
    "compactionStatus" : {
      "message" : "skipped compaction because collection has no datafiles",
      "time" : "2017-04-18T15:52:37Z"
    },
    "readCache" : {
      "count" : 1,
      "size" : 48192,
      "total" : 13107200
    },
    "documentReferences" : 0,
    "waitingFor" : "-",
    "alive" : {
      "count" : 1,
      "size" : 72
    },
    "dead" : {
      "count" : 0,
      "size" : 0,
      "deletion" : 0
    },
    "datafiles" : {
      "count" : 0,
      "fileSize" : 0
    },
    "journals" : {
      "count" : 1,
      "fileSize" : 33554432
    },
    "compactors" : {
      "count" : 0,
      "fileSize" : 0
    },
    "revisions" : {
      "count" : 1,
      "size" : 48192
    }
  },
  "status" : 3,
  "type" : 2,
  "error" : false,
  "code" : 200
}
```

Return Codes

- *200*: Returns information about the collection:

Reply Body

- **count**: The number of documents currently present in the collection.
- **journalSize**: The maximal size of a journal or datafile in bytes.
- **figures**:
  - **datafiles**:
  - **count**: The number of datafiles.
  - **fileSize**: The total filesize of datafiles (in bytes).
  - **uncollectedLogfileEntries**: The number of markers in the write-ahead log for this collection that have not been transferred to journals or datafiles.
  - **lastTick**: The tick of the last marker that was stored in a journal of the collection. This might be 0 if the collection does not yet have a journal.
  - **compactionStatus**:
  - **message**: The action that was performed when the compaction was last run for the collection. This information can be used for debugging compaction issues.
  - **time**: The point in time the compaction for the collection was last executed. This information can be used for debugging compaction issues.
  - **dead**:
  - **count**: The number of dead documents. This includes document versions that have been deleted or replaced by a newer version. Documents deleted or replaced that are contained the write-ahead log only are not reported in this figure.
  - **deletion**: The total number of deletion markers. Deletion markers only contained in the write-ahead log are not reporting in this figure.
  - **size**: The total size in bytes used by all dead documents.
  - **compactors**:
  - **count**: The number of compactor files.
  - **fileSize**: The total filesize of all compactor files (in bytes).
  - **readcache**:
  - **count**: The number of revisions of this collection stored in the document revisions cache.
  - **size**: The memory used for storing the revisions of this collection in the document revisions cache (in bytes). This figure does not include the document data but only mappings from document revision ids to cache entry locations.
  - **waitingFor**: An optional string value that contains information about which object type is at the head of the collection's cleanup queue. This information can be used for debugging compaction and unload issues.
  - **alive**:
  - **count**: The number of currently active documents in all datafiles and journals of the collection. Documents that are contained in the write-ahead log only are not reported in this figure.
  - **size**: The total size in bytes used by all active documents of the collection. Documents that are contained in the write-ahead log only are not reported in this figure.
  - **documentReferences**: The number of references to documents in datafiles that JavaScript code currently holds. This information can be used for debugging compaction and unload issues.
  - **indexes**:
  - **count**: The total number of indexes defined for the collection, including the pre-defined indexes (e.g. primary index).
  - **size**: The total memory allocated for indexes in bytes.
  - **journals**:
  - **count**: The number of journal files.
  - **fileSize**: The total filesize of all journal files (in bytes).
  - **revisions**:
  - **count**: The number of revisions of this collection managed by the storage engine.
  - **size**: The memory used for storing the revisions of this collection in the storage engine (in bytes). This figure does not include the document data but only mappings from document revision ids to storage engine datafile positions.

```
shell> curl --dump - http://localhost:8529/_api/collection/products/figures

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
location: /_db/_system/_api/collection/products/figures
```

show response body

Return collection revision id

Retrieve the collections revision id

`GET /_api/collection/{collection-name}/revision`

Path Parameters

- *collection-name* (required): The name of the collection.

In addition to the above, the result will also contain the collection's revision id. The revision id is a server-generated string that clients can use to check whether data in a collection has changed since the last revision check.

- *revision*: The collection revision id as a string.

Example:

Retrieving the revision of a collection

```
shell> curl --dump - http://localhost:8529/_api/collection/products/revision

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "id" : "9480",
  "name" : "products",
  "isSystem" : false,
  "status" : 3,
  "type" : 2,
  "revision" : "0",
  "error" : false,
  "code" : 200
}
```

Return Codes

- *400*: If the *collection-name* is missing, then a *HTTP 400* is returned.
- *404*: If the *collection-name* is unknown, then a *HTTP 404* is returned.

Examples

Retrieving the revision of a collection

```
shell> curl --dump - http://localhost:8529/_api/collection/products/revision

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Return checksum for the collection

returns a checksum for the specified collection

`GET /_api/collection/{collection-name}/checksum`

Path Parameters

- *collection-name* (required): The name of the collection.

Query Parameters

- *withRevisions* (optional): Whether or not to include document revision ids in the checksum calculation.
- *withData* (optional): Whether or not to include document body data in the checksum calculation.

Will calculate a checksum of the meta-data (keys and optionally revision ids) and optionally the document data in the collection.

The checksum can be used to compare if two collections on different ArangoDB instances contain the same contents. The current revision of the collection is returned too so one can make sure the checksums are calculated for the same state of data.

By default, the checksum will only be calculated on the *_key* system attribute of the documents contained in the collection. For edge collections, the system attributes *_from* and *_to* will also be included in the calculation.

By setting the optional query parameter *withRevisions* to *true*, then revision ids (*_rev* system attributes) are included in the checksumming.

By providing the optional query parameter *withData* with a value of *true*, the user-defined document attributes will be included in the calculation too. **Note**: Including user-defined attributes will make the checksumming slower.

The response is a JSON object with the following attributes:

- *checksum*: The calculated checksum as a number.
- *revision*: The collection revision id as a string.

Note: this method is not available in a cluster.

Example:

Retrieving the checksum of a collection:

```
shell> curl --dump - http://localhost:8529/_api/collection/products/checksum

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "id" : "9125",
  "name" : "products",
  "isSystem" : false,
  "status" : 3,
  "type" : 2,
  "checksum" : "4115365245092230908",
  "revision" : "_U2FB4jy---",
  "error" : false,
  "code" : 200
}
```

Example:

Retrieving the checksum of a collection including the collection data, but not the revisions:

```
shell> curl --dump - http://localhost:8529/_api/collection/products/checksum?withRevisions=false&withData=true

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "id" : "9134",
  "name" : "products",
  "isSystem" : false,
  "status" : 3,
  "type" : 2,
  "checksum" : "7360612632771340455",
  "revision" : "_U2FB4rK---",
  "error" : false,
  "code" : 200
}
```

Return Codes

- *400*: If the *collection-name* is missing, then a *HTTP 400* is returned.
- *404*: If the *collection-name* is unknown, then a *HTTP 404* is returned.

Examples

Retrieving the checksum of a collection:

```
shell> curl --dump - http://localhost:8529/_api/collection/products/checksum

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body
Retrieving the checksum of a collection including the collection data, but not the revisions:

```
shell> curl --dump - http://localhost:8529/_api/collection/products/checksum?
withRevisions=false&withData=true

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

reads all collections

returns all collections

`GET /_api/collection`

Query Parameters

- *excludeSystem* (optional): Whether or not system collections should be excluded from the result.

Returns an object with an attribute *collections* containing an array of all collection descriptions. The same information is also available in the *names* as an object with the collection names as keys.

By providing the optional query parameter *excludeSystem* with a value of *true*, all system collections will be excluded from the response.

Example:

Return information about all collections:

```
shell> curl --dump - http://localhost:8529/_api/collection

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "result" : [
    {
      "id" : "60",
      "name" : "_apps",
      "isSystem" : true,
      "status" : 3,
      "type" : 2
    },
    {
      "id" : "36",
      "name" : "_aqlfunctions",
```

```
      "isSystem" : true,
      "status" : 3,
      "type" : 2
    },
    {
      "id" : "54",
      "name" : "_frontend",
      "isSystem" : true,
      "status" : 3,
      "type" : 2
    },
    {
      "id" : "3",
      "name" : "_graphs",
      "isSystem" : true,
      "status" : 3,
      "type" : 2
    },
    {
      "id" : "58",
      "name" : "_jobs",
      "isSystem" : true,
      "status" : 3,
      "type" : 2
    },
    {
      "id" : "20",
      "name" : "_modules",
      "isSystem" : true,
      "status" : 3,
      "type" : 2
    },
    {
      "id" : "56",
      "name" : "_queues",
      "isSystem" : true,
      "status" : 3,
      "type" : 2
    },
    {
      "id" : "22",
      "name" : "_routing",
      "isSystem" : true,
      "status" : 3,
      "type" : 2
    },
    {
      "id" : "43",
      "name" : "_statistics",
      "isSystem" : true,
      "status" : 3,
      "type" : 2
    },
    {
      "id" : "48",
      "name" : "_statistics15",
      "isSystem" : true,
      "status" : 3,
      "type" : 2
    },
    {
      "id" : "38",
      "name" : "_statisticsRaw",
      "isSystem" : true,
      "status" : 3,
      "type" : 2
    },
    {
      "id" : "6",
      "name" : "_users",
      "isSystem" : true,
      "status" : 3,
      "type" : 2
    },
    {
      "id" : "98",
```

```
      "name" : "animals",
      "isSystem" : false,
      "status" : 3,
      "type" : 2
    },
    {
      "id" : "93",
      "name" : "demo",
      "isSystem" : false,
      "status" : 3,
      "type" : 2
    }
  ],
  "error" : false,
  "code" : 200
}
```

Return Codes

- *200*: The list of collections

Examples

Return information about all collections:

```
shell> curl --dump - http://localhost:8529/_api/collection

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Modifying a Collection

Load collection

loads a collection

`PUT /_api/collection/{collection-name}/load`

Path Parameters

- *collection-name* (required): The name of the collection.

Loads a collection into memory. Returns the collection on success.

The request body object might optionally contain the following attribute:

- *count*: If set, this controls whether the return value should include the number of documents in the collection. Setting *count* to *false* may speed up loading a collection. The default value for *count* is *true*.

On success an object with the following attributes is returned:

- *id*: The identifier of the collection.
- *name*: The name of the collection.
- *count*: The number of documents inside the collection. This is only returned if the *count* input parameters is set to *true* or has not been specified.
- *status*: The status of the collection as number.
- *type*: The collection type. Valid types are:
  - 2: document collection
  - 3: edges collection
- *isSystem*: If *true* then the collection is a system collection.

Example:

```
shell> curl -X PUT --dump - http://localhost:8529/_api/collection/products/load

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "id" : "9484",
  "name" : "products",
  "isSystem" : false,
  "count" : 0,
  "status" : 3,
  "type" : 2,
  "error" : false,
  "code" : 200
}
```

Return Codes

- *400*: If the *collection-name* is missing, then a *HTTP 400* is returned.
- *404*: If the *collection-name* is unknown, then a *HTTP 404* is returned.

Examples

```
shell> curl -X PUT --dump - http://localhost:8529/_api/collection/products/load

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Unload collection

unloads a collection

`PUT /_api/collection/{collection-name}/unload`

Path Parameters

- *collection-name* (required):

Removes a collection from memory. This call does not delete any documents. You can use the collection afterwards; in which case it will be loaded into memory, again. On success an object with the following attributes is returned:

- *id*: The identifier of the collection.
- *name*: The name of the collection.
- *status*: The status of the collection as number.
- *type*: The collection type. Valid types are:
  - 2: document collection
  - 3: edges collection
- *isSystem*: If *true* then the collection is a system collection.

Example:

```
shell> curl -X PUT --dump - http://localhost:8529/_api/collection/products/unload

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "id" : "9504",
  "name" : "products",
  "isSystem" : false,
  "status" : 4,
  "type" : 2,
  "error" : false,
  "code" : 200
}
```

Return Codes

- *400*: If the *collection-name* is missing, then a *HTTP 400* is returned.
- *404*: If the *collection-name* is unknown, then a *HTTP 404* is returned.

Examples

```
shell> curl -X PUT --dump - http://localhost:8529/_api/collection/products/unload

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Change properties of a collection

changes a collection

`PUT /_api/collection/{collection-name}/properties`

Path Parameters

- *collection-name* (required): The name of the collection.

Changes the properties of a collection. Expects an object with the attribute(s)

- *waitForSync*: If *true* then creating or changing a document will wait until the data has been synchronized to disk.
- *journalSize*: The maximal size of a journal or datafile in bytes. The value must be at least `1048576` (1 MB). Note that when changing the journalSize value, it will only have an effect for additional journals or datafiles that are created. Already existing journals or datafiles will not be affected.

On success an object with the following attributes is returned:

- *id*: The identifier of the collection.
- *name*: The name of the collection.
- *waitForSync*: The new value.
- *journalSize*: The new value.
- *status*: The status of the collection as number.
- *type*: The collection type. Valid types are:
  - 2: document collection
  - 3: edges collection
- *isSystem*: If *true* then the collection is a system collection.
- *isVolatile*: If *true* then the collection data will be kept in memory only and ArangoDB will not write or sync the data to disk.
- *doCompact*: Whether or not the collection will be compacted.
- *keyOptions*: JSON object which contains key generation options:
  - *type*: specifies the type of the key generator. The currently available generators are *traditional* and *autoincrement*.
  - *allowUserKeys*: if set to *true*, then it is allowed to supply own key values in the *_key* attribute of a document. If set to *false*, then the key generator is solely responsible for generating keys and supplying own key values in the *_key* attribute of documents is considered an error.

Note: except for *waitForSync*, *journalSize* and *name*, collection properties **cannot be changed** once a collection is created. To rename a collection, the rename endpoint must be used.

Example:

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/collection/products/properties <<EOF
{
  "waitForSync" : true
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "id" : "9489",
  "name" : "products",
  "isSystem" : false,
  "doCompact" : true,
  "isVolatile" : false,
  "journalSize" : 33554432,
  "keyOptions" : {
    "type" : "traditional",
    "allowUserKeys" : true
  },
  "waitForSync" : true,
  "indexBuckets" : 8,
  "status" : 3,
  "type" : 2,
  "error" : false,
  "code" : 200
}
```

Return Codes

- *400*: If the *collection-name* is missing, then a *HTTP 400* is returned.
- *404*: If the *collection-name* is unknown, then a *HTTP 404* is returned.

Examples

```
shell> curl -X PUT --data-binary @- --dump -
http://localhost:8529/_api/collection/products/properties <<EOF
{
  "waitForSync" : true
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Rename collection

renames a collection

`PUT /_api/collection/{collection-name}/rename`

Path Parameters

- *collection-name* (required): The name of the collection to rename.

Renames a collection. Expects an object with the attribute(s)

- *name*: The new name.

It returns an object with the attributes

- *id*: The identifier of the collection.
- *name*: The new name of the collection.

- *status*: The status of the collection as number.
- *type*: The collection type. Valid types are:
  - 2: document collection
  - 3: edges collection
- *isSystem*: If *true* then the collection is a system collection.

If renaming the collection succeeds, then the collection is also renamed in all graph definitions inside the `_graphs` collection in the current database.

Note: this method is not available in a cluster.

Example:

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/collection/products1/rename <<EOF
{
  "name" : "newname"
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "id" : "9495",
  "name" : "newname",
  "isSystem" : false,
  "status" : 3,
  "type" : 2,
  "error" : false,
  "code" : 200
}
```

Return Codes

- *400*: If the *collection-name* is missing, then a *HTTP 400* is returned.
- *404*: If the *collection-name* is unknown, then a *HTTP 404* is returned.

Examples

```
shell> curl -X PUT --data-binary @- --dump -
http://localhost:8529/_api/collection/products1/rename <<EOF
{
  "name" : "newname"
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Rotate journal of a collection

rotates the journal of a collection

`PUT /_api/collection/{collection-name}/rotate`

Path Parameters

- *collection-name* (required): The name of the collection.

Rotates the journal of a collection. The current journal of the collection will be closed and made a read-only datafile. The purpose of the rotate method is to make the data in the file available for compaction (compaction is only performed for read-only datafiles, and not for journals).

Saving new data in the collection subsequently will create a new journal file automatically if there is no current journal.

It returns an object with the attributes

- *result*: will be *true* if rotation succeeded

Note: This method is not available in a cluster.

Example:

Rotating the journal:

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/collection/products/rotate <<EOF
{
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "result" : true,
  "error" : false,
  "code" : 200
}
```

Example:

Rotating if no journal exists:

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/collection/products/rotate <<EOF
{
}
EOF

HTTP/1.1 400 Bad Request
content-type: application/json; charset=utf-8

{
  "error" : true,
  "code" : 400,
  "errorNum" : 1105,
  "errorMessage" : "could not rotate journal: no journal"
}
```

Return Codes

- *400*: If the collection currently has no journal, *HTTP 400* is returned.
- *404*: If the *collection-name* is unknown, then a *HTTP 404* is returned.

Examples

Rotating the journal:

```
shell> curl -X PUT --data-binary @- --dump -
http://localhost:8529/_api/collection/products/rotate <<EOF
{
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Rotating if no journal exists:

```
shell> curl -X PUT --data-binary @- --dump -
http://localhost:8529/_api/collection/products/rotate <<EOF
{
}
EOF

HTTP/1.1 400 Bad Request
content-type: application/json; charset=utf-8
```

show response body

HTTP Interface for Documents

In this chapter we describe the REST API of ArangoDB for documents.

- Basic approach
- Detailed API description

Basics and Terminology

Documents, Keys, Handles and Revisions

Documents in ArangoDB are JSON objects. These objects can be nested (to any depth) and may contain lists. Each document has a unique primary key which identifies it within its collection. Furthermore, each document is uniquely identified by its document handle across all collections in the same database. Different revisions of the same document (identified by its handle) can be distinguished by their document revision. Any transaction only ever sees a single revision of a document.

Here is an example document:

```
{
  "_id" : "myusers/3456789",
  "_key" : "3456789",
  "_rev" : "14253647",
  "firstName" : "John",
  "lastName" : "Doe",
  "address" : {
    "street" : "Road To Nowhere 1",
    "city" : "Gotham"
  },
  "hobbies" : [
    {name: "swimming", howFavorite: 10},
    {name: "biking", howFavorite: 6},
    {name: "programming", howFavorite: 4}
  ]
}
```

All documents contain special attributes: the document handle is stored as a string in `_id`, the document's primary key in `_key` and the document revision in `_rev`. The value of the `_key` attribute can be specified by the user when creating a document. `_id` and `_key` values are immutable once the document has been created. The `_rev` value is maintained by ArangoDB automatically.

Document Handle

A document handle uniquely identifies a document in the database. It is a string and consists of the collection's name and the document key (`_key` attribute) separated by `/`.

Document Key

A document key uniquely identifies a document in the collection it is stored in. It can and should be used by clients when specific documents are queried. The document key is stored in the `_key` attribute of each document. The key values are automatically indexed by ArangoDB in a collection's primary index. Thus looking up a document by its key is a fast operation. The _key value of a document is immutable once the document has been created. By default, ArangoDB will auto-generate a document key if no _key attribute is specified, and use the user-specified _key otherwise.

This behavior can be changed on a per-collection level by creating collections with the `keyOptions` attribute.

Using `keyOptions` it is possible to disallow user-specified keys completely, or to force a specific regime for auto-generating the `_key` values.

Document Revision

As ArangoDB supports MVCC (Multiple Version Concurrency Control), documents can exist in more than one revision. The document revision is the MVCC token used to specify a particular revision of a document (identified by its `_id`). It is a string value currently containing an integer number and is unique within the list of document revisions for a single document. Document revisions can be used to conditionally query, update, replace or delete documents in the database. In order to find a particular revision of a document, you need the document handle or key, and the document revision.

ArangoDB uses 64bit unsigned integer values to maintain document revisions internally. When returning document revisions to clients, ArangoDB will put them into a string to ensure the revision is not clipped by clients that do not support big integers. Clients should treat the revision returned by ArangoDB as an opaque string when they store or use it locally. This will allow ArangoDB to change the format of revisions later if this should be required. Clients can use revisions to perform simple equality/non-equality comparisons (e.g. to check whether a document has changed or not), but they should not use revision ids to perform greater/less than comparisons with them to check if a document revision is older than one another, even if this might work for some cases.

Document Etag

ArangoDB tries to adhere to the existing HTTP standard as far as possible. To this end, results of single document queries have the HTTP header `ETag` set to the document revision enclosed in double quotes.

The basic operations (create, read, exists, replace, update, delete) for documents are mapped to the standard HTTP methods (*POST*, *GET*, *HEAD*, *PUT*, *PATCH* and *DELETE*).

If you modify a document, you can use the *If-Match* field to detect conflicts. The revision of a document can be checking using the HTTP method *HEAD*.

Multiple Documents in a single Request

Beginning with ArangoDB 3.0 the basic document API has been extended to handle not only single documents but multiple documents in a single request. This is crucial for performance, in particular in the cluster situation, in which a single request can involve multiple network hops within the cluster. Another advantage is that it reduces the overhead of the HTTP protocol and individual network round trips between the client and the server. The general idea to perform multiple document operations in a single request is to use a JSON array of objects in the place of a single document. As a consequence, document keys, handles and revisions for preconditions have to be supplied embedded in the individual documents given. Multiple document operations are restricted to a single document or edge collections. See the API descriptions for details.

Note that the *GET*, *HEAD* and *DELETE* HTTP operations generally do not allow to pass a message body. Thus, they cannot be used to perform multiple document operations in one request. However, there are other endpoints to request and delete multiple documents in one request. FIXME: ADD SENSIBLE LINKS HERE.

URI of a Document

Any document can be retrieved using its unique URI:

```
http://server:port/_api/document/<document-handle>
```

For example, assuming that the document handle is `demo/362549736`, then the URL of that document is:

```
http://localhost:8529/_api/document/demo/362549736
```

The above URL schema does not specify a database name explicitly, so the default database `_system` will be used. To explicitly specify the database context, use the following URL schema:

```
http://server:port/_db/<database-name>/_api/document/<document-handle>
```

Example:

```
http://localhost:8529/_db/mydb/_api/document/demo/362549736
```

Note: The following examples use the short URL format for brevity.

The document revision is returned in the "ETag" HTTP header when requesting a document.

If you obtain a document using *GET* and you want to check whether a newer revision is available, then you can use the *If-None-Match* header. If the document is unchanged, a *HTTP 412* (precondition failed) error is returned.

If you want to query, replace, update or delete a document, then you can use the *If-Match* header. If the document has changed, then the operation is aborted and an *HTTP 412* error is returned.

Working with Documents using REST

Read document

reads a single document

`GET /_api/document/{document-handle}`

Path Parameters

- *document-handle* (required): The handle of the document.

Header Parameters

- *If-None-Match* (optional): If the "If-None-Match" header is given, then it must contain exactly one ETag. The document is returned, if it has a different revision than the given ETag. Otherwise an *HTTP 304* is returned.
- *If-Match* (optional): If the "If-Match" header is given, then it must contain exactly one ETag. The document is returned, if it has the same revision as the given ETag. Otherwise a *HTTP 412* is returned.

Returns the document identified by *document-handle*. The returned document contains three special attributes: *_id* containing the document handle, *_key* containing key which uniquely identifies a document in a given collection and *_rev* containing the revision.

Example:

Use a document handle:

```
shell> curl --dump - http://localhost:8529/_api/document/products/9968

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
etag: "_U2FCGca---"

{
  "_key" : "9968",
  "_id" : "products/9968",
  "_rev" : "_U2FCGca---",
  "hello" : "world"
}
```

Example:

Use a document handle and an etag:

```
shell> curl --header 'If-None-Match: "_U2FCG42---"' --dump - http://localhost:8529/_api/document/products/10012
```

Example:

Unknown document handle:

```
shell> curl --dump - http://localhost:8529/_api/document/products/unknownhandle

HTTP/1.1 404 Not Found
content-type: application/json; charset=utf-8

{
  "error" : true,
  "errorMessage" : "collection 'products' not found",
  "code" : 404,
  "errorNum" : 1203
}
```

Return Codes

- *200*: is returned if the document was found
- *304*: is returned if the "If-None-Match" header is given and the document has the same version
- *404*: is returned if the document or collection was not found
- *412*: is returned if an "If-Match" header is given and the found document has a different version. The response will also contain the found document's current revision in the *_rev* attribute. Additionally, the attributes *_id* and *_key* will be returned.

Examples

Use a document handle:

```
shell> curl --dump - http://localhost:8529/_api/document/products/9968

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
etag: "_U2FCGca---"
```

show response body
Use a document handle and an etag:

```
shell> curl --header 'If-None-Match: "_U2FCG42---"' --dump -
http://localhost:8529/_api/document/products/10012
```

Unknown document handle:

```
shell> curl --dump - http://localhost:8529/_api/document/products/unknownhandle

HTTP/1.1 404 Not Found
content-type: application/json; charset=utf-8
```

show response body

Changes in 3.0 from 2.8:

The *rev* query parameter has been withdrawn. The same effect can be achieved with the *If-Match* HTTP header.

Read document header

reads a single document head

`HEAD /_api/document/{document-handle}`

Path Parameters

- *document-handle* (required): The handle of the document.

Header Parameters

- *If-None-Match* (optional): If the "If-None-Match" header is given, then it must contain exactly one ETag. If the current document revision is not equal to the specified ETag, an *HTTP 200* response is returned. If the current document revision is identical to the specified ETag, then an *HTTP 304* is returned.
- *If-Match* (optional): If the "If-Match" header is given, then it must contain exactly one ETag. The document is returned, if it has the same revision as the given ETag. Otherwise a *HTTP 412* is returned.

Like *GET*, but only returns the header fields and not the body. You can use this call to get the current revision of a document or check if the document was deleted.

Example:

```
shell> curl -X HEAD --dump - http://localhost:8529/_api/document/products/10004
```

Return Codes

- *200*: is returned if the document was found
- *304*: is returned if the "If-None-Match" header is given and the document has the same version
- *404*: is returned if the document or collection was not found
- *412*: is returned if an "If-Match" header is given and the found document has a different version. The response will also contain the found document's current revision in the *ETag* header.

Examples

```
shell> curl -X HEAD --dump - http://localhost:8529/_api/document/products/10004
```

Changes in 3.0 from 2.8:

The *rev* query parameter has been withdrawn. The same effect can be achieved with the *If-Match* HTTP header.

Read all documents

reads all documents from collection

`PUT /_api/simple/all-keys`

Query Parameters

A json post document with these Properties is required:

- **type**: The type of the result. The following values are allowed:
  - *id*: returns an array of document ids (*_id* attributes)
  - *key*: returns an array of document keys (*_key* attributes)
  - *path*: returns an array of document URI paths. This is the default.
- **collection**: The name of the collection. This is only for backward compatibility. In ArangoDB versions < 3.0, the URL path was */_api/document* and this was passed in via the query parameter "collection". This combination was removed.

Returns an array of all keys, ids, or URI paths for all documents in the collection identified by *collection*. The type of the result array is determined by the *type* attribute.

Note that the results have no defined order and thus the order should not be relied on.

Example:

Return all document paths

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/all-keys <<EOF
{
  "collection" : "products"
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8

{
  "result" : [
    "/_db/_system/_api/document/products/9994",
    "/_db/_system/_api/document/products/9997",
    "/_db/_system/_api/document/products/9990"
  ],
  "hasMore" : false,
  "cached" : false,
  "extra" : {
    "stats" : {
      "writesExecuted" : 0,
      "writesIgnored" : 0,
      "scannedFull" : 3,
      "scannedIndex" : 0,
      "filtered" : 0,
      "executionTime" : 0.0003190040588378906
    },
    "warnings" : [ ]
  },
  "error" : false,
  "code" : 201
}
```

Example:

Return all document keys

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/all-keys <<EOF
{
  "collection" : "products",
  "type" : "id"
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8

{
  "result" : [
    "products/9983",
    "products/9980",
    "products/9976"
  ],
  "hasMore" : false,
  "cached" : false,
  "extra" : {
    "stats" : {
      "writesExecuted" : 0,
      "writesIgnored" : 0,
      "scannedFull" : 3,
      "scannedIndex" : 0,
      "filtered" : 0,
      "executionTime" : 0.0003330707550048828
    },
    "warnings" : [ ]
  },
  "error" : false,
  "code" : 201
}
```

Example:

Collection does not exist

```
shell> curl --dump - http://localhost:8529/_api/document/doesnotexist

HTTP/1.1 404 Not Found
content-type: application/json; charset=utf-8

{
  "error" : true,
  "errorMessage" : "expecting GET /_api/document/<document-handle>",
  "code" : 404,
  "errorNum" : 1203
}
```

Return Codes

- *201*: All went well.
- *404*: The collection does not exist.

Examples

Return all document paths

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/all-keys
<<EOF
{
  "collection" : "products"
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8
```

show response body

Return all document keys

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/all-keys
<<EOF
{
  "collection" : "products",
  "type" : "id"
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8
```

show response body

Collection does not exist

```
shell> curl --dump - http://localhost:8529/_api/document/doesnotexist

HTTP/1.1 404 Not Found
content-type: application/json; charset=utf-8
```

show response body

Changes in 3.0 from 2.8:

The collection name should now be specified in the URL path. The old way with the URL path */_api/document* and the required query parameter *collection* still works.

Create document

creates documents

`POST /_api/document/{collection}`

Path Parameters

- *collection* (required): The *collection* in which the collection is to be created.

Post Body

data (required): A JSON representation of a single document or of an array of documents.

Query Parameters

- *collection* (optional): The name of the collection. This is only for backward compatibility. In ArangoDB versions < 3.0, the URL path was */_api/document* and this query parameter was required. This combination still works, but the recommended way is to specify the collection in the URL path.
- *waitForSync* (optional): Wait until document has been synced to disk.
- *returnNew* (optional): Additionally return the complete new document under the attribute *new* in the result.
- *silent* (optional): If set to *true*, an empty object will be returned as response. No meta-data will be returned for the created document. This option can be used to save some network traffic.

free style json body

Creates a new document from the document given in the body, unless there is already a document with the *_key* given. If no *_key* is given, a new unique *_key* is generated automatically.

The body can be an array of documents, in which case all documents in the array are inserted with the same semantics as for a single document. The result body will contain a JSON array of the same length as the input array, and each entry contains the result of the operation for the corresponding input. In case of an error the entry is a document with attributes *error* set to *true* and errorCode set to the error code that has happened.

Possibly given *_id* and *_rev* attributes in the body are always ignored, the URL part or the query parameter collection respectively counts.

If the document was created successfully, then the *Location* header contains the path to the newly created document. The *ETag* header field contains the revision of the document. Both are only set in the single document case.

If *silent* is not set to *true*, the body of the response contains a JSON object (single document case) with the following attributes:

- *_id* contains the document handle of the newly created document
- *_key* contains the document key
- *_rev* contains the document revision

In the multi case the body is an array of such objects.

If the collection parameter *waitForSync* is *false*, then the call returns as soon as the document has been accepted. It will not wait until the documents have been synced to disk.

Optionally, the query parameter *waitForSync* can be used to force synchronization of the document creation operation to disk even in case that the *waitForSync* flag had been disabled for the entire collection. Thus, the *waitForSync* query parameter can be used to force synchronization of just this specific operations. To use this, set the *waitForSync* parameter to *true*. If the *waitForSync* parameter is not specified or set to *false*, then the collection's default *waitForSync* behavior is applied. The *waitForSync* query parameter cannot be used to disable synchronization for collections that have a default *waitForSync* value of *true*.

If the query parameter *returnNew* is *true*, then, for each generated document, the complete new document is returned under the *new* attribute in the result.

Example:

Create a document in a collection named *products*. Note that the revision identifier might or might not by equal to the auto-generated key.

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/document/products <<EOF
{ "Hello": "World" }
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8
etag: "_U2FCGAe---"
location: /_db/_system/_api/document/products/9934

{
  "_id" : "products/9934",
  "_key" : "9934",
  "_rev" : "_U2FCGAe---"
}
```

Example:

Create a document in a collection named *products* with a collection-level *waitForSync* value of *false*.

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/document/products <<EOF
{ "Hello": "World" }
EOF

HTTP/1.1 202 Accepted
content-type: application/json; charset=utf-8
etag: "_U2FCFxK---"
location: /_db/_system/_api/document/products/9924

{
  "_id" : "products/9924",
  "_key" : "9924",
  "_rev" : "_U2FCFxK---"
}
```

Example:

Create a document in a collection with a collection-level *waitForSync* value of *false*, but using the *waitForSync* query parameter.

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/document/products?waitForSync=true <<EOF
{ "Hello": "World" }
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8
etag: "_U2FCGV----"
location: /_db/_system/_api/document/products/9961

{
  "_id" : "products/9961",
  "_key" : "9961",
  "_rev" : "_U2FCGV----"
}
```

Example:

Unknown collection name

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/document/products <<EOF
{ "Hello": "World" }
EOF

HTTP/1.1 404 Not Found
content-type: application/json; charset=utf-8

{
  "error" : true,
  "errorMessage" : "collection 'products' not found",
  "code" : 404,
  "errorNum" : 1203
}
```

Example:

Illegal document

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/document/products <<EOF
{ 1: "World" }
EOF

HTTP/1.1 400 Bad Request
content-type: application/json; charset=utf-8

{
  "error" : true,
  "errorMessage" : "VPackError error: Expecting '\"' or '}'",
  "code" : 400,
  "errorNum" : 600
}
```

Example:

Insert multiple documents:

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/document/products <<EOF
[{"Hello":"Earth"}, {"Hello":"Venus"}, {"Hello":"Mars"}]
EOF

HTTP/1.1 202 Accepted
content-type: application/json; charset=utf-8

[
  {
    "_id" : "products/9941",
    "_key" : "9941",
    "_rev" : "_U2FCGH2---"
  },
  {
    "_id" : "products/9945",
    "_key" : "9945",
    "_rev" : "_U2FCGH2--_"
  },
  {
    "_id" : "products/9947",
    "_key" : "9947",
    "_rev" : "_U2FCGH2--A"
  }
]
```

Example:

Use of returnNew:

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/document/products?returnNew=true <<EOF
{"Hello":"World"}
EOF

HTTP/1.1 202 Accepted
content-type: application/json; charset=utf-8
etag: "_U2FCGNq---"
location: /_db/_system/_api/document/products/9954

{
  "_id" : "products/9954",
  "_key" : "9954",
  "_rev" : "_U2FCGNq---",
  "new" : {
    "_key" : "9954",
    "_id" : "products/9954",
    "_rev" : "_U2FCGNq---",
    "Hello" : "World"
  }
}
```

Return Codes

- *201*: is returned if the documents were created successfully and *waitForSync* was *true*.
- *202*: is returned if the documents were created successfully and *waitForSync* was *false*.
- *400*: is returned if the body does not contain a valid JSON representation of one document or an array of documents. The response body contains an error document in this case.
- *404*: is returned if the collection specified by *collection* is unknown. The response body contains an error document in this case.
- *409*: is returned in the single document case if a document with the same qualifiers in an indexed attribute conflicts with an already existing document and thus violates that unique constraint. The response body contains an error document in this case. In the array case only 201 or 202 is returned, but if an error occurred, the additional HTTP header *X-Arango-Error-Codes* is set, which contains a map of the error codes that occurred together with their multiplicities, as in: *1205:10,1210:17* which means that in 10 cases the error 1205 "illegal document handle" and in 17 cases the error 1210 "unique constraint violated" has happened.

Examples

Create a document in a collection named *products*. Note that the revision identifier might or might not by equal to the auto-generated key.

```
shell> curl -X POST --data-binary @- --dump -
http://localhost:8529/_api/document/products <<EOF
{ "Hello": "World" }
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8
etag: "_U2FCGAe---"
location: /_db/_system/_api/document/products/9934
```

show response body

Create a document in a collection named *products* with a collection-level *waitForSync* value of *false*.

```
shell> curl -X POST --data-binary @- --dump -
http://localhost:8529/_api/document/products <<EOF
{ "Hello": "World" }
EOF

HTTP/1.1 202 Accepted
content-type: application/json; charset=utf-8
etag: "_U2FCFxK---"
location: /_db/_system/_api/document/products/9924
```

show response body

Create a document in a collection with a collection-level *waitForSync* value of *false*, but using the *waitForSync* query parameter.

```
shell> curl -X POST --data-binary @- --dump -
http://localhost:8529/_api/document/products?waitForSync=true <<EOF
{ "Hello": "World" }
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8
etag: "_U2FCGV----"
location: /_db/_system/_api/document/products/9961
```

show response body

Unknown collection name

```
shell> curl -X POST --data-binary @- --dump -
http://localhost:8529/_api/document/products <<EOF
{ "Hello": "World" }
EOF

HTTP/1.1 404 Not Found
content-type: application/json; charset=utf-8
```

show response body

Illegal document

```
shell> curl -X POST --data-binary @- --dump -
http://localhost:8529/_api/document/products <<EOF
{ 1: "World" }
EOF

HTTP/1.1 400 Bad Request
content-type: application/json; charset=utf-8
```

show response body

Insert multiple documents:

```
shell> curl -X POST --data-binary @- --dump -
http://localhost:8529/_api/document/products <<EOF
[{"Hello":"Earth"}, {"Hello":"Venus"}, {"Hello":"Mars"}]
EOF

HTTP/1.1 202 Accepted
content-type: application/json; charset=utf-8

[
  {
    "_id" : "products/9941",
    "_key" : "9941",
    "_rev" : "_U2FCGH2---"
  },
  {
    "_id" : "products/9945",
    "_key" : "9945",
    "_rev" : "_U2FCGH2--_"
  },
  {
    "_id" : "products/9947",
    "_key" : "9947",
    "_rev" : "_U2FCGH2--A"
  }
]
```

Use of returnNew:

```
shell> curl -X POST --data-binary @- --dump -
http://localhost:8529/_api/document/products?returnNew=true <<EOF
{"Hello":"World"}
EOF

HTTP/1.1 202 Accepted
content-type: application/json; charset=utf-8
etag: "_U2FCGNq---"
location: /_db/_system/_api/document/products/9954
```

show response body

Changes in 3.0 from 2.8:

The collection name should now be specified in the URL path. The old way with the URL path */_api/document* and the required query parameter *collection* still works. The possibility to insert multiple documents with one operation is new and the query parameter *returnNew* has been added.

Replace document

replaces a document

`PUT /_api/document/{document-handle}`

Post Body

document (required): A JSON representation of a single document.

Path Parameters

- *document-handle* (required): This URL parameter must be a document handle.

Query Parameters

- *waitForSync* (optional): Wait until document has been synced to disk.
- *ignoreRevs* (optional): By default, or if this is set to *true*, the *_rev* attributes in the given document is ignored. If this is set to *false*, then the *_rev* attribute given in the body document is taken as a precondition. The document is only replaced if the current revision is the one specified.
- *returnOld* (optional): Return additionally the complete previous revision of the changed document under the attribute *old* in the result.
- *returnNew* (optional): Return additionally the complete new document under the attribute *new* in the result.
- *silent* (optional): If set to *true*, an empty object will be returned as response. No meta-data will be returned for the replaced document. This option can be used to save some network traffic.

Header Parameters

- *If-Match* (optional): You can conditionally replace a document based on a target revision id by using the *if-match* HTTP header.

free style json body

Replaces the document with handle with the one in the body, provided there is such a document and no precondition is violated.

If the *If-Match* header is specified and the revision of the document in the database is unequal to the given revision, the precondition is violated.

If *If-Match* is not given and *ignoreRevs* is *false* and there is a *_rev* attribute in the body and its value does not match the revision of the document in the database, the precondition is violated.

If a precondition is violated, an *HTTP 412* is returned.

If the document exists and can be updated, then an *HTTP 201* or an *HTTP 202* is returned (depending on *waitForSync*, see below), the *ETag* header field contains the new revision of the document and the *Location* header contains a complete URL under which the document can be queried.

Optionally, the query parameter *waitForSync* can be used to force synchronization of the document replacement operation to disk even in case that the *waitForSync* flag had been disabled for the entire collection. Thus, the *waitForSync* query parameter can be used to force synchronization of just specific operations. To use this, set the *waitForSync* parameter to *true*. If the *waitForSync* parameter is not specified or set to *false*, then the collection's default *waitForSync* behavior is applied. The *waitForSync* query parameter cannot be used to disable synchronization for collections that have a default *waitForSync* value of *true*.

If *silent* is not set to *true*, the body of the response contains a JSON object with the information about the handle and the revision. The attribute *_id* contains the known *document-handle* of the updated document, *_key* contains the key which uniquely identifies a document in a given collection, and the attribute *_rev* contains the new document revision.

If the query parameter *returnOld* is *true*, then the complete previous revision of the document is returned under the *old* attribute in the result.

If the query parameter *returnNew* is *true*, then the complete new document is returned under the *new* attribute in the result.

If the document does not exist, then a *HTTP 404* is returned and the body of the response contains an error document.

Example:

Using a document handle

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/document/products/10020 <<EOF
{"Hello": "you"}
EOF

HTTP/1.1 202 Accepted
content-type: application/json; charset=utf-8
etag: "_U2FCH_y---"
location: /_db/_system/_api/document/products/10020

{
  "_id" : "products/10020",
  "_key" : "10020",
  "_rev" : "_U2FCH_y---",
  "_oldRev" : "_U2FCH_u---"
}
```

Example:

Unknown document handle

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/document/products/10040 <<EOF
{}
EOF

HTTP/1.1 404 Not Found
content-type: application/json; charset=utf-8

{
  "error" : true,
  "errorMessage" : "document not found",
  "code" : 404,
  "errorNum" : 1202
}
```

Example:

Produce a revision conflict

```
shell> curl -X PUT --header 'If-Match: "_U2FCHGm---"' --data-binary @- --dump - http://localhost:8529/_api/document/p
roducts/10029 <<EOF
{"other":"content"}
EOF

HTTP/1.1 412 Precondition Failed
content-type: application/json; charset=utf-8
etag: "_U2FCHGi---"

{
  "error" : true,
  "code" : 412,
  "errorNum" : 1200,
  "errorMessage" : "precondition failed",
  "_id" : "products/10029",
  "_key" : "10029",
  "_rev" : "_U2FCHGi---"
}
```

Return Codes

- *201*: is returned if the document was replaced successfully and *waitForSync* was *true*.
- *202*: is returned if the document was replaced successfully and *waitForSync* was *false*.

- *400*: is returned if the body does not contain a valid JSON representation of a document. The response body contains an error document in this case.
- *404*: is returned if the collection or the document was not found.
- *412*: is returned if the precondition was violated. The response will also contain the found documents' current revisions in the *_rev* attributes. Additionally, the attributes *_id* and *_key* will be returned.

Examples

Using a document handle

```
shell> curl -X PUT --data-binary @- --dump -
http://localhost:8529/_api/document/products/10020 <<EOF
{"Hello": "you"}
EOF

HTTP/1.1 202 Accepted
content-type: application/json; charset=utf-8
etag: "_U2FCH_y---"
location: /_db/_system/_api/document/products/10020
```

show response body

Unknown document handle

```
shell> curl -X PUT --data-binary @- --dump -
http://localhost:8529/_api/document/products/10040 <<EOF
{}
EOF

HTTP/1.1 404 Not Found
content-type: application/json; charset=utf-8
```

show response body

Produce a revision conflict

```
shell> curl -X PUT --header 'If-Match: "_U2FCHGm---"' --data-binary @- --dump -
http://localhost:8529/_api/document/products/10029 <<EOF
{"other":"content"}
EOF

HTTP/1.1 412 Precondition Failed
content-type: application/json; charset=utf-8
etag: "_U2FCHGi---"
```

show response body

Changes in 3.0 from 2.8:

There are quite some changes in this in comparison to Version 2.8, but few break existing usage:

- the *rev* query parameter is gone (was duplication of If-Match)
- the *policy* query parameter is gone (was non-sensical)
- the *ignoreRevs* query parameter is new, the default *true* gives the traditional behavior as in 2.8
- the *returnNew* and *returnOld* query parameters are new

There should be very few changes to behavior happening in real-world situations or drivers. Essentially, one has to replace usage of the *rev* query parameter by usage of the *If-Match* header. The non-sensical combination of *If-Match* given and *policy=last* no longer works, but can easily be achieved by leaving out the *If-Match* header.

The collection name should now be specified in the URL path. The old way with the URL path */_api/document* and the required query parameter *collection* still works.

Replace documents

replaces multiple documents

`PUT /_api/document/{collection}`

Post Body

documents (required): A JSON representation of an array of documents.

Path Parameters

- *collection* (required): This URL parameter is the name of the collection in which the documents are replaced.

Query Parameters

- *waitForSync* (optional): Wait until the new documents have been synced to disk.
- *ignoreRevs* (optional): By default, or if this is set to *true*, the *_rev* attributes in the given documents are ignored. If this is set to *false*, then any *_rev* attribute given in a body document is taken as a precondition. The document is only replaced if the current revision is the one specified.
- *returnOld* (optional): Return additionally the complete previous revision of the changed documents under the attribute *old* in the result.
- *returnNew* (optional): Return additionally the complete new documents under the attribute *new* in the result.

free style json body

Replaces multiple documents in the specified collection with the ones in the body, the replaced documents are specified by the *_key* attributes in the body documents.

If *ignoreRevs* is *false* and there is a *_rev* attribute in a document in the body and its value does not match the revision of the corresponding document in the database, the precondition is violated.

If the document exists and can be updated, then an *HTTP 201* or an *HTTP 202* is returned (depending on *waitForSync*, see below).

Optionally, the query parameter *waitForSync* can be used to force synchronization of the document replacement operation to disk even in case that the *waitForSync* flag had been disabled for the entire collection. Thus, the *waitForSync* query parameter can be used to force synchronization of just specific operations. To use this, set the *waitForSync* parameter to *true*. If the *waitForSync* parameter is not specified or set to *false*, then the collection's default *waitForSync* behavior is applied. The *waitForSync* query parameter cannot be used to disable synchronization for collections that have a default *waitForSync* value of *true*.

The body of the response contains a JSON array of the same length as the input array with the information about the handle and the revision of the replaced documents. In each entry, the attribute *_id* contains the known *document-handle* of each updated document, *_key* contains the key which uniquely identifies a document in a given collection, and the attribute *_rev* contains the new document revision. In case of an error or violated precondition, an error object with the attribute *error* set to *true* and the attribute *errorCode* set to the error code is built.

If the query parameter *returnOld* is *true*, then, for each generated document, the complete previous revision of the document is returned under the *old* attribute in the result.

If the query parameter *returnNew* is *true*, then, for each generated document, the complete new document is returned under the *new* attribute in the result.

Note that if any precondition is violated or an error occurred with some of the documents, the return code is still 201 or 202, but the additional HTTP header *X-Arango-Error-Codes* is set, which contains a map of the error codes that occurred together with their multiplicities, as in: *1200:17,1205:10* which means that in 17 cases the error 1200 "revision conflict" and in 10 cases the error 1205 "illegal document handle" has happened.

Return Codes

- *201*: is returned if the documents were replaced successfully and *waitForSync* was *true*.
- *202*: is returned if the documents were replaced successfully and *waitForSync* was *false*.
- *400*: is returned if the body does not contain a valid JSON representation of an array of documents. The response body contains an error document in this case.
- *404*: is returned if the collection was not found.

Changes in 3.0 from 2.8:

The multi document version is new in 3.0.

Update document

updates a document

`PATCH /_api/document/{document-handle}`

Post Body

document (required): A JSON representation of a document update as an object.

Path Parameters

- *document-handle* (required): This URL parameter must be a document handle.

Query Parameters

- *keepNull* (optional): If the intention is to delete existing attributes with the patch command, the URL query parameter *keepNull* can be used with a value of *false*. This will modify the behavior of the patch command to remove any attributes from the existing document that are contained in the patch document with an attribute value of *null*.
- *mergeObjects* (optional): Controls whether objects (not arrays) will be merged if present in both the existing and the patch document. If set to *false*, the value in the patch document will overwrite the existing document's value. If set to *true*, objects will be merged. The default is *true*.
- *waitForSync* (optional): Wait until document has been synced to disk.
- *ignoreRevs* (optional): By default, or if this is set to *true*, the *_rev* attributes in the given document is ignored. If this is set to *false*, then the *_rev* attribute given in the body document is taken as a precondition. The document is only updated if the current revision is the one specified.
- *returnOld* (optional): Return additionally the complete previous revision of the changed document under the attribute *old* in the result.
- *returnNew* (optional): Return additionally the complete new document under the attribute *new* in the result.
- *silent* (optional): If set to *true*, an empty object will be returned as response. No meta-data will be returned for the updated document. This option can be used to save some network traffic.

Header Parameters

- *If-Match* (optional): You can conditionally update a document based on a target revision id by using the *if-match* HTTP header.

free style json body

Partially updates the document identified by *document-handle*. The body of the request must contain a JSON document with the attributes to patch (the patch document). All attributes from the patch document will be added to the existing document if they do not yet exist, and overwritten in the existing document if they do exist there.

Setting an attribute value to *null* in the patch document will cause a value of *null* to be saved for the attribute by default.

If the *If-Match* header is specified and the revision of the document in the database is unequal to the given revision, the precondition is violated.

If *If-Match* is not given and *ignoreRevs* is *false* and there is a *_rev* attribute in the body and its value does not match the revision of the document in the database, the precondition is violated.

If a precondition is violated, an *HTTP 412* is returned.

If the document exists and can be updated, then an *HTTP 201* or an *HTTP 202* is returned (depending on *waitForSync*, see below), the *ETag* header field contains the new revision of the document and the *Location* header contains a complete URL under which the document can be queried.

Optionally, the query parameter *waitForSync* can be used to force synchronization of the updated document operation to disk even in case that the *waitForSync* flag had been disabled for the entire collection. Thus, the *waitForSync* query parameter can be used to force synchronization of just specific operations. To use this, set the *waitForSync* parameter to *true*. If the *waitForSync* parameter is not specified or set to *false*, then the collection's default *waitForSync* behavior is applied. The *waitForSync* query parameter cannot be used to disable synchronization for collections that have a default *waitForSync* value of *true*.

If *silent* is not set to *true*, the body of the response contains a JSON object with the information about the handle and the revision. The attribute *_id* contains the known *document-handle* of the updated document, *_key* contains the key which uniquely identifies a document in a given collection, and the attribute *_rev* contains the new document revision.

If the query parameter *returnOld* is *true*, then the complete previous revision of the document is returned under the *old* attribute in the result.

If the query parameter *returnNew* is *true*, then the complete new document is returned under the *new* attribute in the result.

If the document does not exist, then a *HTTP 404* is returned and the body of the response contains an error document.

Example:

Patches an existing document with new content.

```
shell> curl -X PATCH --data-binary @- --dump - http://localhost:8529/_api/document/products/9895 <<EOF
{
  "hello" : "world"
}
EOF

HTTP/1.1 202 Accepted
content-type: application/json; charset=utf-8
etag: "_U2FCFg---_"
location: /_db/_system/_api/document/products/9895

{
  "_id" : "products/9895",
  "_key" : "9895",
  "_rev" : "_U2FCFg---_",
  "_oldRev" : "_U2FCFg----"
}
shell> curl -X PATCH --data-binary @- --dump - http://localhost:8529/_api/document/products/9895 <<EOF
{
  "numbers" : {
    "one" : 1,
    "two" : 2,
    "three" : 3,
    "empty" : null
  }
}
EOF

HTTP/1.1 202 Accepted
content-type: application/json; charset=utf-8
etag: "_U2FCFgO---"
location: /_db/_system/_api/document/products/9895

{
  "_id" : "products/9895",
  "_key" : "9895",
  "_rev" : "_U2FCFgO---",
  "_oldRev" : "_U2FCFg---_"
}
shell> curl --dump - http://localhost:8529/_api/document/products/9895

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
etag: "_U2FCFgO---"
```

```
{
  "_key" : "9895",
  "_id" : "products/9895",
  "_rev" : "_U2FCFgO---",
  "one" : "world",
  "hello" : "world",
  "numbers" : {
    "one" : 1,
    "two" : 2,
    "three" : 3,
    "empty" : null
  }
}
shell> curl -X PATCH --data-binary @- --dump - http://localhost:8529/_api/document/products/9895?keepNull=false <<EOF
{
  "hello" : null,
  "numbers" : {
    "four" : 4
  }
}
EOF

HTTP/1.1 202 Accepted
content-type: application/json; charset=utf-8
etag: "_U2FCFgq---"
location: /_db/_system/_api/document/products/9895

{
  "_id" : "products/9895",
  "_key" : "9895",
  "_rev" : "_U2FCFgq---",
  "_oldRev" : "_U2FCFgO---"
}
shell> curl --dump - http://localhost:8529/_api/document/products/9895

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
etag: "_U2FCFgq---"

{
  "_key" : "9895",
  "_id" : "products/9895",
  "_rev" : "_U2FCFgq---",
  "one" : "world",
  "numbers" : {
    "empty" : null,
    "one" : 1,
    "three" : 3,
    "two" : 2,
    "four" : 4
  }
}
```

Example:

Merging attributes of an object using `mergeObjects`:

```
shell> curl --dump - http://localhost:8529/_api/document/products/9910

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
etag: "_U2FCFo2---"

{
  "_key" : "9910",
  "_id" : "products/9910",
  "_rev" : "_U2FCFo2---",
  "inhabitants" : {
    "china" : 1366980000,
    "india" : 1263590000,
    "usa" : 319220000
```

```
  }
}
shell> curl -X PATCH --data-binary @- --dump - http://localhost:8529/_api/document/products/9910?mergeObjects=true <<
EOF
{
  "inhabitants" : {
    "indonesia" : 252164800,
    "brazil" : 203553000
  }
}
EOF

shell> curl --dump - http://localhost:8529/_api/document/products/9910

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
etag: "_U2FCFpC---"

{
  "_key" : "9910",
  "_id" : "products/9910",
  "_rev" : "_U2FCFpC---",
  "inhabitants" : {
    "china" : 1366980000,
    "india" : 1263590000,
    "usa" : 319220000,
    "indonesia" : 252164800,
    "brazil" : 203553000
  }
}
shell> curl -X PATCH --data-binary @- --dump - http://localhost:8529/_api/document/products/9910?mergeObjects=false <
<EOF
{
  "inhabitants" : {
    "pakistan" : 188346000
  }
}
EOF

HTTP/1.1 202 Accepted
content-type: application/json; charset=utf-8
etag: "_U2FCFpW---"
location: /_db/_system/_api/document/products/9910

{
  "_id" : "products/9910",
  "_key" : "9910",
  "_rev" : "_U2FCFpW---",
  "_oldRev" : "_U2FCFpC---"
}
shell> curl --dump - http://localhost:8529/_api/document/products/9910

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
etag: "_U2FCFpW---"

{
  "_key" : "9910",
  "_id" : "products/9910",
  "_rev" : "_U2FCFpW---",
  "inhabitants" : {
    "pakistan" : 188346000
  }
}
```

Return Codes

- *201*: is returned if the document was updated successfully and *waitForSync* was *true*.
- *202*: is returned if the document was updated successfully and *waitForSync* was *false*.
- *400*: is returned if the body does not contain a valid JSON representation of a document. The response body contains an error document in this case.
- *404*: is returned if the collection or the document was not found.
- *412*: is returned if the precondition was violated. The response will also contain the found documents' current revisions in the *_rev*

attributes. Additionally, the attributes *_id* and *_key* will be returned.

Examples

Patches an existing document with new content.

```
shell> curl -X PATCH --data-binary @- --dump -
http://localhost:8529/_api/document/products/9895 <<EOF
{
  "hello" : "world"
}
EOF

HTTP/1.1 202 Accepted
content-type: application/json; charset=utf-8
etag: "_U2FCFg---_"
location: /_db/_system/_api/document/products/9895
```

show response body

Merging attributes of an object using `mergeObjects`:

```
shell> curl --dump - http://localhost:8529/_api/document/products/9910

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
etag: "_U2FCFo2---"
```

show response body

Changes in 3.0 from 2.8:

There are quite some changes in this in comparison to Version 2.8, but few break existing usage:

- the *rev* query parameter is gone (was duplication of If-Match)
- the *policy* query parameter is gone (was non-sensical)
- the *ignoreRevs* query parameter is new, the default *true* gives the traditional behavior as in 2.8
- the *returnNew* and *returnOld* query parameters are new

There should be very few changes to behavior happening in real-world situations or drivers. Essentially, one has to replace usage of the *rev* query parameter by usage of the *If-Match* header. The non-sensical combination of *If-Match* given and *policy=last* no longer works, but can easily be achieved by leaving out the *If-Match* header.

The collection name should now be specified in the URL path. The old way with the URL path */_api/document* and the required query parameter *collection* still works.

Update documents

updates multiple documents

`PATCH /_api/document/{collection}`

Post Body

documents (required): A JSON representation of an array of document updates as objects.

Path Parameters

- *collection* (required): This URL parameter is the name of the collection in which the documents are updated.

Query Parameters

- *keepNull* (optional): If the intention is to delete existing attributes with the patch command, the URL query parameter *keepNull* can be used with a value of *false*. This will modify the behavior of the patch command to remove any attributes from the existing document that are contained in the patch document with an attribute value of *null*.
- *mergeObjects* (optional): Controls whether objects (not arrays) will be merged if present in both the existing and the patch document. If set to *false*, the value in the patch document will overwrite the existing document's value. If set to *true*, objects will be merged. The default is *true*.
- *waitForSync* (optional): Wait until the new documents have been synced to disk.
- *ignoreRevs* (optional): By default, or if this is set to *true*, the *_rev* attributes in the given documents are ignored. If this is set to *false*, then any *_rev* attribute given in a body document is taken as a precondition. The document is only updated if the current revision is the one specified.
- *returnOld* (optional): Return additionally the complete previous revision of the changed documents under the attribute *old* in the result.
- *returnNew* (optional): Return additionally the complete new documents under the attribute *new* in the result.

free style json body

Partially updates documents, the documents to update are specified by the *_key* attributes in the body objects. The body of the request must contain a JSON array of document updates with the attributes to patch (the patch documents). All attributes from the patch documents will be added to the existing documents if they do not yet exist, and overwritten in the existing documents if they do exist there.

Setting an attribute value to *null* in the patch documents will cause a value of *null* to be saved for the attribute by default.

If *ignoreRevs* is *false* and there is a *_rev* attribute in a document in the body and its value does not match the revision of the corresponding document in the database, the precondition is violated.

If the document exists and can be updated, then an *HTTP 201* or an *HTTP 202* is returned (depending on *waitForSync*, see below).

Optionally, the query parameter *waitForSync* can be used to force synchronization of the document replacement operation to disk even in case that the *waitForSync* flag had been disabled for the entire collection. Thus, the *waitForSync* query parameter can be used to force synchronization of just specific operations. To use this, set the *waitForSync* parameter to *true*. If the *waitForSync* parameter is not specified or set to *false*, then the collection's default *waitForSync* behavior is applied. The *waitForSync* query parameter cannot be used to disable synchronization for collections that have a default *waitForSync* value of *true*.

The body of the response contains a JSON array of the same length as the input array with the information about the handle and the revision of the updated documents. In each entry, the attribute *_id* contains the known *document-handle* of each updated document, *_key* contains the key which uniquely identifies a document in a given collection, and the attribute *_rev* contains the new document revision. In case of an error or violated precondition, an error object with the attribute *error* set to *true* and the attribute *errorCode* set to the error code is built.

If the query parameter *returnOld* is *true*, then, for each generated document, the complete previous revision of the document is returned under the *old* attribute in the result.

If the query parameter *returnNew* is *true*, then, for each generated document, the complete new document is returned under the *new* attribute in the result.

Note that if any precondition is violated or an error occurred with some of the documents, the return code is still 201 or 202, but the additional HTTP header *X-Arango-Error-Codes* is set, which contains a map of the error codes that occurred together with their multiplicities, as in: *1200:17,1205:10* which means that in 17 cases the error 1200 "revision conflict" and in 10 cases the error 1205 "illegal document handle" has happened.

Return Codes

- *201*: is returned if the documents were updated successfully and *waitForSync* was *true*.
- *202*: is returned if the documents were updated successfully and *waitForSync* was *false*.
- *400*: is returned if the body does not contain a valid JSON representation of an array of documents. The response body contains an error document in this case.
- *404*: is returned if the collection was not found.

Changes in 3.0 from 2.8:

The multi document version is new in 3.0.

Removes a document

removes a document

`DELETE /_api/document/{document-handle}`

Path Parameters

- *document-handle* (required): Removes the document identified by *document-handle*.

Query Parameters

- *waitForSync* (optional): Wait until deletion operation has been synced to disk.
- *returnOld* (optional): Return additionally the complete previous revision of the changed document under the attribute *old* in the result.
- *silent* (optional): If set to *true*, an empty object will be returned as response. No meta-data will be returned for the removed document. This option can be used to save some network traffic.

Header Parameters

- *If-Match* (optional): You can conditionally remove a document based on a target revision id by using the *if-match* HTTP header.

If *silent* is not set to *true*, the body of the response contains a JSON object with the information about the handle and the revision. The attribute *_id* contains the known *document-handle* of the removed document, *_key* contains the key which uniquely identifies a document in a given collection, and the attribute *_rev* contains the document revision.

If the *waitForSync* parameter is not specified or set to *false*, then the collection's default *waitForSync* behavior is applied. The *waitForSync* query parameter cannot be used to disable synchronization for collections that have a default *waitForSync* value of *true*.

If the query parameter *returnOld* is *true*, then the complete previous revision of the document is returned under the *old* attribute in the result.

Example:

Using document handle:

```
shell> curl -X DELETE --dump - http://localhost:8529/_api/document/products/9835

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
etag: "_U2FCEo2---"
location: /_db/_system/_api/document/products/9835

{
  "_id" : "products/9835",
  "_key" : "9835",
  "_rev" : "_U2FCEo2---"
}
```

Example:

Unknown document handle:

```
shell> curl -X DELETE --dump - http://localhost:8529/_api/document/products/9875

HTTP/1.1 404 Not Found
content-type: application/json; charset=utf-8

{
  "error" : true,
  "errorMessage" : "document not found",
  "code" : 404,
  "errorNum" : 1202
}
```

Example:

Revision conflict:

```
shell> curl -X DELETE --header 'If-Match: "_U2FCE0a---"' --dump - http://localhost:8529/_api/document/products/9844

HTTP/1.1 412 Precondition Failed
content-type: application/json; charset=utf-8
etag: "_U2FCE0W---"

{
  "error" : true,
  "code" : 412,
  "errorNum" : 1200,
  "errorMessage" : "precondition failed",
  "_id" : "products/9844",
  "_key" : "9844",
  "_rev" : "_U2FCE0W---"
}
```

Return Codes

- *200*: is returned if the document was removed successfully and *waitForSync* was *true*.
- *202*: is returned if the document was removed successfully and *waitForSync* was *false*.
- *404*: is returned if the collection or the document was not found. The response body contains an error document in this case.
- *412*: is returned if a "If-Match" header or *rev* is given and the found document has a different version. The response will also contain the found document's current revision in the *_rev* attribute. Additionally, the attributes *_id* and *_key* will be returned.

Examples

Using document handle:

```
shell> curl -X DELETE --dump - http://localhost:8529/_api/document/products/9835

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
etag: "_U2FCEo2---"
location: /_db/_system/_api/document/products/9835
```

show response body
Unknown document handle:

```
shell> curl -X DELETE --dump - http://localhost:8529/_api/document/products/9875

HTTP/1.1 404 Not Found
content-type: application/json; charset=utf-8
```

show response body
Revision conflict:

```
shell> curl -X DELETE --header 'If-Match: "_U2FCE0a---"' --dump -
http://localhost:8529/_api/document/products/9844

HTTP/1.1 412 Precondition Failed
content-type: application/json; charset=utf-8
etag: "_U2FCE0W---"
```

show response body

Changes in 3.0 from 2.8:

There are only very few changes in this in comparison to Version 2.8:

- the *rev* query parameter is gone (was duplication of If-Match)
- the *policy* query parameter is gone (was non-sensical)
- the *returnOld* query parameter is new

There should be very few changes to behavior happening in real-world situations or drivers. Essentially, one has to replace usage of the *rev* query parameter by usage of the *If-Match* header. The non-sensical combination of *If-Match* given and *policy=last* no longer works, but can easily be achieved by leaving out the *If-Match* header.

Removes multiple documents

removes multiple document

`DELETE /_api/document/{collection}`

Post Body

array (required): A JSON array of strings or documents.

Path Parameters

- *collection* (required): Collection from which documents are removed.

Query Parameters

- *waitForSync* (optional): Wait until deletion operation has been synced to disk.
- *returnOld* (optional): Return additionally the complete previous revision of the changed document under the attribute *old* in the result.
- *ignoreRevs* (optional): If set to *true*, ignore any *_rev* attribute in the selectors. No revision check is performed.

free style json body

The body of the request is an array consisting of selectors for documents. A selector can either be a string with a key or a string with a document handle or an object with a *_key* attribute. This API call removes all specified documents from *collection*. If the selector is an object and has a *_rev* attribute, it is a precondition that the actual revision of the removed document in the collection is the specified one.

The body of the response is an array of the same length as the input array. For each input selector, the output contains a JSON object with the information about the outcome of the operation. If no error occurred, an object is built in which the attribute *_id* contains the known *document-handle* of the removed document, *_key* contains the key which uniquely identifies a document in a given collection, and the attribute *_rev* contains the document revision. In case of an error, an object with the attribute *error* set to *true* and *errorCode* set to the error code is built.

If the *waitForSync* parameter is not specified or set to *false*, then the collection's default *waitForSync* behavior is applied. The *waitForSync* query parameter cannot be used to disable synchronization for collections that have a default *waitForSync* value of *true*.

If the query parameter *returnOld* is *true*, then the complete previous revision of the document is returned under the *old* attribute in the result.

Note that if any precondition is violated or an error occurred with some of the documents, the return code is still 200 or 202, but the additional HTTP header *X-Arango-Error-Codes* is set, which contains a map of the error codes that occurred together with their multiplicities, as in: *1200:17,1205:10* which means that in 17 cases the error 1200 "revision conflict" and in 10 cases the error 1205 "illegal document handle" has happened.

Example:

Using document handle:

```
shell> curl -X DELETE --dump - http://localhost:8529/_api/document/products/9866

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
etag: "_U2FCFCC---"
location: /_db/_system/_api/document/products/9866

{
  "_id" : "products/9866",
  "_key" : "9866",
  "_rev" : "_U2FCFCC---"
}
```

Example:

Unknown document handle:

```
shell> curl -X DELETE --dump - http://localhost:8529/_api/document/products/9885

HTTP/1.1 404 Not Found
content-type: application/json; charset=utf-8

{
  "error" : true,
  "errorMessage" : "document not found",
  "code" : 404,
  "errorNum" : 1202
}
```

Example:

Revision conflict:

```
shell> curl -X DELETE --header 'If-Match: "_U2FCE8C--_"' --dump - http://localhost:8529/_api/document/products/9855

HTTP/1.1 412 Precondition Failed
content-type: application/json; charset=utf-8
etag: "_U2FCE8C---"

{
  "error" : true,
  "code" : 412,
  "errorNum" : 1200,
  "errorMessage" : "precondition failed",
  "_id" : "products/9855",
  "_key" : "9855",
  "_rev" : "_U2FCE8C---"
}
```

Return Codes

- *200*: is returned if *waitForSync* was *true*.
- *202*: is returned if *waitForSync* was *false*.

- *404*: is returned if the collection was not found. The response body contains an error document in this case.

Examples

Using document handle:

```
shell> curl -X DELETE --dump - http://localhost:8529/_api/document/products/9866

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
etag: "_U2FCFCC---"
location: /_db/_system/_api/document/products/9866
```

show response body
Unknown document handle:

```
shell> curl -X DELETE --dump - http://localhost:8529/_api/document/products/9885

HTTP/1.1 404 Not Found
content-type: application/json; charset=utf-8
```

show response body
Revision conflict:

```
shell> curl -X DELETE --header 'If-Match: "_U2FCE8C--_"' --dump -
http://localhost:8529/_api/document/products/9855

HTTP/1.1 412 Precondition Failed
content-type: application/json; charset=utf-8
etag: "_U2FCE8C---"
```

show response body

Changes in 3.0 from 2.8:

This variant is new in 3.0. Note that it requires a body in the DELETE request.

HTTP Interface for Edges

This is an introduction to ArangoDB's REST interface for edges.

ArangoDB offers graph functionality; Edges are one part of that.

Address and ETag of an Edge

All documents in ArangoDB have a document handle. This handle uniquely identifies a document. Any document can be retrieved using its unique URI:

```
http://server:port/_api/document/<document-handle>
```

Edges are a special variation of documents. To access an edge use the same URL format as for a document:

```
http://server:port/_api/document/<document-handle>
```

For example, assumed that the document handle, which is stored in the *_id* attribute of the edge, is *demo/362549736*, then the URL of that edge is:

```
http://localhost:8529/_api/document/demo/362549736
```

The above URL scheme does not specify a database name explicitly, so the default database will be used. To explicitly specify the database context, use the following URL schema:

```
http://server:port/_db/<database-name>/_api/document/<document-handle>
```

Example:

```
http://localhost:8529/_db/mydb/_api/document/demo/362549736
```

Note: that the following examples use the short URL format for brevity.

Working with Edges using REST

This is documentation to ArangoDB's REST interface for edges.

Edges are documents with two additional attributes: *_from* and *_to*. These attributes are mandatory and must contain the document-handle of the from and to vertices of an edge.

Use the general document REST api for create/read/update/delete.

Read in- or outbound edges

get edges

`GET /_api/edges/{collection-id}`

Path Parameters

- *collection-id* (required): The id of the collection.

Query Parameters

- *vertex* (required): The id of the start vertex.
- *direction* (optional): Selects *in* or *out* direction for edges. If not set, any edges are returned.

Returns an array of edges starting or ending in the vertex identified by *vertex-handle*.

Example:

Any direction

```
shell> curl --dump - http://localhost:8529/_api/edges/edges?vertex=vertices/1

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "edges" : [
    {
      "_key" : "5",
      "_id" : "edges/5",
      "_from" : "vertices/1",
      "_to" : "vertices/3",
      "_rev" : "_U2FCHYC---",
      "$label" : "v1 -> v3"
    },
    {
      "_key" : "6",
      "_id" : "edges/6",
      "_from" : "vertices/2",
      "_to" : "vertices/1",
      "_rev" : "_U2FCHYG---",
      "$label" : "v2 -> v1"
    },
    {
      "_key" : "7",
      "_id" : "edges/7",
      "_from" : "vertices/4",
      "_to" : "vertices/1",
      "_rev" : "_U2FCHYG--_",
      "$label" : "v4 -> v1"
    }
  ],
  "error" : false,
  "code" : 200,
  "stats" : {
    "scannedIndex" : 3,
    "filtered" : 0
  }
}
```

Example:

In edges

```
shell> curl --dump - http://localhost:8529/_api/edges/edges?vertex=vertices/1&direction=in

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "edges" : [
    {
      "_key" : "6",
      "_id" : "edges/6",
      "_from" : "vertices/2",
      "_to" : "vertices/1",
      "_rev" : "_U2FCHpG--_",
      "$label" : "v2 -> v1"
    },
    {
      "_key" : "7",
      "_id" : "edges/7",
      "_from" : "vertices/4",
      "_to" : "vertices/1",
      "_rev" : "_U2FCHpK---",
      "$label" : "v4 -> v1"
    }
  ],
  "error" : false,
  "code" : 200,
  "stats" : {
    "scannedIndex" : 2,
    "filtered" : 0
  }
}
```

Example:

Out edges

```
shell> curl --dump - http://localhost:8529/_api/edges/edges?vertex=vertices/1&direction=out

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "edges" : [
    {
      "_key" : "5",
      "_id" : "edges/5",
      "_from" : "vertices/1",
      "_to" : "vertices/3",
      "_rev" : "_U2FCH8u--_",
      "$label" : "v1 -> v3"
    }
  ],
  "error" : false,
  "code" : 200,
  "stats" : {
    "scannedIndex" : 1,
    "filtered" : 0
  }
}
```

Return Codes

- *200*: is returned if the edge collection was found and edges were retrieved.
- *400*: is returned if the request contains invalid parameters.
- *404*: is returned if the edge collection was not found.

Examples

Any direction

```
shell> curl --dump - http://localhost:8529/_api/edges/edges?vertex=vertices/1

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

In edges

```
shell> curl --dump - http://localhost:8529/_api/edges/edges?
vertex=vertices/1&direction=in

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Out edges

```
shell> curl --dump - http://localhost:8529/_api/edges/edges?
vertex=vertices/1&direction=out

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

General Graphs

This chapter describes the REST interface for the multi-collection graph module. It allows you to define a graph that is spread across several edge and document collections. There is no need to include the referenced collections within the query, this module will handle it for you.

Manage your graphs

The graph module provides functions dealing with graph structures. Examples will explain the REST API on the social graph:

List all graphs

Lists all graphs known to the graph module.

`GET /_api/gharial`

Lists all graph names stored in this database.

Example:

```
shell> curl --dump - http://localhost:8529/_api/gharial

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "error" : false,
  "graphs" : [
    {
      "_key" : "routeplanner",
      "_id" : "_graphs/routeplanner",
      "_rev" : "_U2FBV4e---",
      "orphanCollections" : [ ],
      "edgeDefinitions" : [
        {
          "collection" : "germanHighway",
          "from" : [
            "germanCity"
          ],
          "to" : [
            "germanCity"
          ]
        },
        {
          "collection" : "frenchHighway",
          "from" : [
            "frenchCity"
          ],
          "to" : [
            "frenchCity"
          ]
        },
        {
          "collection" : "internationalHighway",
          "from" : [
            "frenchCity",
            "germanCity"
          ],
          "to" : [
            "frenchCity",
            "germanCity"
          ]
        }
      ]
    },
    {
      "_key" : "social",
      "_id" : "_graphs/social",
      "_rev" : "_U2FBVoi---",
      "orphanCollections" : [ ],
      "edgeDefinitions" : [
        {
          "collection" : "relation",
          "from" : [
            "female",
            "male"
          ],
          "to" : [
            "female",
            "male"
          ]
        }
      ]
    }
  ],
  "code" : 200
}
```

Return Codes

- *200*: Is returned if the module is available and the graphs could be listed.

Examples

```
shell> curl --dump - http://localhost:8529/_api/gharial

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Create a graph

Create a new graph in the graph module.

`POST /_api/gharial`

The creation of a graph requires the name of the graph and a definition of its edges. See also edge definitions.

A json post document with these Properties is required:

- **orphanCollections**: An array of additional vertex collections.
- **edgeDefinitions**: An array of definitions for the edge
- **name**: Name of the graph.
- **isSmart**: Define if the created graph should be smart. This only has effect in Enterprise version.
- **options**:
  - **smartGraphAttribute**: The attribute name that is used to smartly shard the vertices of a graph. Every vertex in this Graph has to have this attribute. Cannot be modified later.
  - **numberOfShards**: The number of shards that is used for every collection within this graph. Cannot be modified later.

Example:

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/gharial <<EOF
{
  "name" : "myGraph",
  "edgeDefinitions" : [
    {
      "collection" : "edges",
      "from" : [
        "startVertices"
      ],
      "to" : [
        "endVertices"
      ]
    }
  ]
}
EOF

HTTP/1.1 202 Accepted
etag: _U2FBSeS---
content-type: application/json; charset=utf-8

{
  "error" : false,
  "graph" : {
    "name" : "myGraph",
    "edgeDefinitions" : [
      {
        "collection" : "edges",
        "from" : [
          "startVertices"
        ],
        "to" : [
          "endVertices"
        ]
      }
    ],
    "orphanCollections" : [ ],
    "isSmart" : false,
    "numberOfShards" : 0,
    "smartGraphAttribute" : "",
    "_id" : "_graphs/myGraph",
    "_rev" : "_U2FBSeS---"
  },
  "code" : 202
}
```

Example:

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/gharial <<EOF
{
  "name" : "myGraph",
  "edgeDefinitions" : [
    {
      "collection" : "edges",
      "from" : [
        "startVertices"
      ],
      "to" : [
        "endVertices"
      ]
    }
  ],
  "isSmart" : true,
  "options" : {
    "numberOfShards" : 9,
    "smartGraphAttribute" : "region"
  }
}
EOF

HTTP/1.1 202 Accepted
etag: _U2FBSxy---
content-type: application/json; charset=utf-8

{
  "error" : false,
  "graph" : {
    "name" : "myGraph",
    "edgeDefinitions" : [
      {
        "collection" : "edges",
        "from" : [
          "startVertices"
        ],
        "to" : [
          "endVertices"
        ]
      }
    ],
    "orphanCollections" : [ ],
    "isSmart" : false,
    "numberOfShards" : 0,
    "smartGraphAttribute" : "",
    "_id" : "_graphs/myGraph",
    "_rev" : "_U2FBSxy---"
  },
  "code" : 202
}
```

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/gharial <<EOF
{
  "name" : "myGraph",
  "edgeDefinitions" : [
    {
      "collection" : "edges",
      "from" : [
        "startVertices"
      ],
      "to" : [
        "endVertices"
      ]
    }
  ]
}
EOF

HTTP/1.1 202 Accepted
etag: _U2FBSeS---
content-type: application/json; charset=utf-8
```

show response body

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/gharial <<EOF
{
  "name" : "myGraph",
  "edgeDefinitions" : [
    {
      "collection" : "edges",
      "from" : [
        "startVertices"
      ],
      "to" : [
        "endVertices"
      ]
    }
  ],
  "isSmart" : true,
  "options" : {
    "numberOfShards" : 9,
    "smartGraphAttribute" : "region"
  }
}
EOF

HTTP/1.1 202 Accepted
etag: _U2FBSxy---
content-type: application/json; charset=utf-8
```

show response body

Get a graph

Get a graph from the graph module.

`GET /_api/gharial/{graph-name}`

Gets a graph from the collection *_graphs*. Returns the definition content of this graph.

Example:

```
shell> curl --dump - http://localhost:8529/_api/gharial/myGraph

HTTP/1.1 200 OK
etag: _U2FBU7a---
content-type: application/json; charset=utf-8

{
  "error" : false,
  "graph" : {
    "name" : "myGraph",
    "edgeDefinitions" : [
      {
        "collection" : "edges",
        "from" : [
          "startVertices"
        ],
        "to" : [
          "endVertices"
        ]
      }
    ],
    "orphanCollections" : [ ],
    "isSmart" : false,
    "numberOfShards" : 0,
    "smartGraphAttribute" : "",
    "_id" : "_graphs/myGraph",
    "_rev" : "_U2FBU7a---"
  },
  "code" : 200
}
```

Path Parameters

- *graph-name* (required): The name of the graph.

Return Codes

- *200*: Returned if the graph could be found.
- *404*: Returned if no graph with this name could be found.

Examples

```
shell> curl --dump - http://localhost:8529/_api/gharial/myGraph

HTTP/1.1 200 OK
etag: _U2FBU7a---
content-type: application/json; charset=utf-8
```

show response body

Drop a graph

delete an existing graph

`DELETE /_api/gharial/{graph-name}`

Removes a graph from the collection *_graphs*.

Example:

```
shell> curl -X DELETE --dump - http://localhost:8529/_api/gharial/social?dropCollections=true

HTTP/1.1 202 Accepted
content-type: application/json; charset=utf-8

{
  "error" : false,
  "removed" : true,
  "code" : 202
}
```

Path Parameters

- *graph-name* (required): The name of the graph.

Query Parameters

- *dropCollections* (optional): Drop collections of this graph as well. Collections will only be dropped if they are not used in other graphs.

Return Codes

- *201*: Is returned if the graph could be dropped and waitForSync is enabled for the `_graphs` collection.
- *202*: Returned if the graph could be dropped and waitForSync is disabled for the `_graphs` collection.
- *404*: Returned if no graph with this name could be found.

Examples

```
shell> curl -X DELETE --dump - http://localhost:8529/_api/gharial/social?
dropCollections=true

HTTP/1.1 202 Accepted
content-type: application/json; charset=utf-8
```

show response body

List vertex collections

Lists all vertex collections used in this graph.

`GET /_api/gharial/{graph-name}/vertex`

Lists all vertex collections within this graph.

Example:

```
shell> curl --dump - http://localhost:8529/_api/gharial/social/vertex

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "error" : false,
  "collections" : [
    "female",
    "male"
  ],
  "code" : 200
}
```

Path Parameters

- *graph-name* (required): The name of the graph.

Return Codes

- *200*: Is returned if the collections could be listed.
- *404*: Returned if no graph with this name could be found.

Examples

```
shell> curl --dump - http://localhost:8529/_api/gharial/social/vertex

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Add vertex collection

Add an additional vertex collection to the graph.

`POST /_api/gharial/{graph-name}/vertex`

Adds a vertex collection to the set of collections of the graph. If the collection does not exist, it will be created.

Example:

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/gharial/social/vertex <<EOF
{
  "collection" : "otherVertices"
}
EOF

HTTP/1.1 202 Accepted
etag: _U2FBRna---
content-type: application/json; charset=utf-8

{
  "error" : false,
  "graph" : {
    "name" : "social",
    "edgeDefinitions" : [
      {
        "collection" : "relation",
        "from" : [
          "female",
          "male"
        ],
        "to" : [
          "female",
          "male"
        ]
      }
    ],
    "orphanCollections" : [
      "otherVertices"
    ],
    "isSmart" : false,
    "numberOfShards" : 0,
    "smartGraphAttribute" : "",
    "_id" : "_graphs/social",
    "_rev" : "_U2FBRna---"
  },
  "code" : 202
}
```

Path Parameters

- *graph-name* (required): The name of the graph.

Return Codes

- *201*: Returned if the edge collection could be added successfully and waitForSync is true.
- *202*: Returned if the edge collection could be added successfully and waitForSync is false.
- *404*: Returned if no graph with this name could be found.

Examples

```
shell> curl -X POST --data-binary @- --dump -
http://localhost:8529/_api/gharial/social/vertex <<EOF
{
  "collection" : "otherVertices"
}
EOF

HTTP/1.1 202 Accepted
etag: _U2FBRna---
content-type: application/json; charset=utf-8
```

show response body

Remove vertex collection

Remove a vertex collection form the graph.

`DELETE /_api/gharial/{graph-name}/vertex/{collection-name}`

Removes a vertex collection from the graph and optionally deletes the collection, if it is not used in any other graph.

Example:

You can remove vertex collections that are not used in any edge collection:

```
shell> curl -X DELETE --dump - http://localhost:8529/_api/gharial/social/vertex/otherVertices

HTTP/1.1 202 Accepted
etag: _U2FBYdK---
content-type: application/json; charset=utf-8

{
  "error" : false,
  "graph" : {
    "name" : "social",
    "edgeDefinitions" : [
      {
        "collection" : "relation",
        "from" : [
          "female",
          "male"
        ],
        "to" : [
          "female",
          "male"
        ]
      }
    ],
    "orphanCollections" : [ ],
    "isSmart" : false,
    "numberOfShards" : 0,
    "smartGraphAttribute" : "",
    "_id" : "_graphs/social",
    "_rev" : "_U2FBYdK---"
  },
  "code" : 202
}
```

Example:

You cannot remove vertex collections that are used in edge collections:

```
shell> curl -X DELETE --dump - http://localhost:8529/_api/gharial/social/vertex/male

HTTP/1.1 400 Bad Request
content-type: application/json

{
  "error" : true,
  "errorNum" : 1928,
  "errorMessage" : "not in orphan collection",
  "code" : 400
}
```

Path Parameters

- *graph-name* (required): The name of the graph.
- *collection-name* (required): The name of the vertex collection.

Query Parameters

- *dropCollection* (optional): Drop the collection as well. Collection will only be dropped if it is not used in other graphs.

Return Codes

- *201*: Returned if the vertex collection was removed from the graph successfully and waitForSync is true.
- *202*: Returned if the request was successful but waitForSync is false.
- *400*: Returned if the vertex collection is still used in an edge definition. In this case it cannot be removed from the graph yet, it has to be removed from the edge definition first.
- *404*: Returned if no graph with this name could be found.

Examples

You can remove vertex collections that are not used in any edge collection:

```
shell> curl -X DELETE --dump -
http://localhost:8529/_api/gharial/social/vertex/otherVertices

HTTP/1.1 202 Accepted
etag: _U2FBYdK---
content-type: application/json; charset=utf-8
```

show response body

You cannot remove vertex collections that are used in edge collections:

```
shell> curl -X DELETE --dump - http://localhost:8529/_api/gharial/social/vertex/male

HTTP/1.1 400 Bad Request
content-type: application/json
```

show response body

List edge definitions

Lists all edge definitions

`GET /_api/gharial/{graph-name}/edge`

Lists all edge collections within this graph.

Example:

```
shell> curl --dump - http://localhost:8529/_api/gharial/social/edge

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "error" : false,
  "collections" : [
    "relation"
  ],
  "code" : 200
}
```

Path Parameters

- *graph-name* (required): The name of the graph.

Return Codes

- *200*: Is returned if the edge definitions could be listed.
- *404*: Returned if no graph with this name could be found.

Examples

```
shell> curl --dump - http://localhost:8529/_api/gharial/social/edge

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Add edge definition

Add a new edge definition to the graph

`POST /_api/gharial/{graph-name}/edge`

Adds an additional edge definition to the graph.

This edge definition has to contain a *collection* and an array of each *from* and *to* vertex collections. An edge definition can only be added if this definition is either not used in any other graph, or it is used with exactly the same definition. It is not possible to store a definition "e" from "v1" to "v2" in the one graph, and "e" from "v2" to "v1" in the other graph.

A json post document with these Properties is required:

- **to** (string): One or many vertex collections that can contain target vertices.
- **from** (string): One or many vertex collections that can contain source vertices.
- **collection**: The name of the edge collection to be used.

Example:

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/gharial/social/edge <<EOF
{
  "collection" : "works_in",
  "from" : [
    "female",
    "male"
  ],
  "to" : [
    "city"
  ]
}
EOF

HTTP/1.1 202 Accepted
etag: _U2FBQwu---
content-type: application/json; charset=utf-8

{
  "error" : false,
  "graph" : {
    "name" : "social",
    "edgeDefinitions" : [
      {
        "collection" : "relation",
        "from" : [
          "female",
          "male"
        ],
        "to" : [
          "female",
          "male"
        ]
      },
      {
        "collection" : "works_in",
        "from" : [
          "female",
          "male"
        ],
        "to" : [
          "city"
        ]
      }
    ],
    "orphanCollections" : [ ],
    "isSmart" : false,
    "numberOfShards" : 0,
    "smartGraphAttribute" : "",
    "_id" : "_graphs/social",
    "_rev" : "_U2FBQwu---"
  },
  "code" : 202
}
```

Path Parameters

- *graph-name* (required): The name of the graph.

Return Codes

- *201*: Returned if the definition could be added successfully and waitForSync is enabled for the `_graphs` collection.
- *202*: Returned if the definition could be added successfully and waitForSync is disabled for the `_graphs` collection.
- *400*: Returned if the defininition could not be added, the edge collection is used in an other graph with a different signature.
- *404*: Returned if no graph with this name could be found.

Examples

```
shell> curl -X POST --data-binary @- --dump -
http://localhost:8529/_api/gharial/social/edge <<EOF
{
  "collection" : "works_in",
  "from" : [
    "female",
    "male"
  ],
  "to" : [
    "city"
  ]
}
EOF

HTTP/1.1 202 Accepted
etag: _U2FBQwu---
content-type: application/json; charset=utf-8
```

show response body

Replace an edge definition

Replace an existing edge definition

`PUT /_api/gharial/{graph-name}/edge/{definition-name}`

Change one specific edge definition. This will modify all occurrences of this definition in all graphs known to your database.

A json post document with these Properties is required:

- **to** (string): One or many vertex collections that can contain target vertices.
- **from** (string): One or many vertex collections that can contain source vertices.
- **collection**: The name of the edge collection to be used.

Example:

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/gharial/social/edge/relation <<EOF
{
  "collection" : "relation",
  "from" : [
    "female",
    "male",
    "animal"
  ],
  "to" : [
    "female",
    "male",
    "animal"
  ]
}
EOF

HTTP/1.1 202 Accepted
etag: _U2FBZKu---
content-type: application/json; charset=utf-8

{
  "error" : false,
  "graph" : {
    "name" : "social",
    "edgeDefinitions" : [
      {
        "collection" : "relation",
        "from" : [
          "animal",
          "female",
          "male"
        ],
        "to" : [
          "animal",
          "female",
          "male"
        ]
      }
    ],
    "orphanCollections" : [ ],
    "isSmart" : false,
    "numberOfShards" : 0,
    "smartGraphAttribute" : "",
    "_id" : "_graphs/social",
    "_rev" : "_U2FBZKu---"
  },
  "code" : 202
}
```

Path Parameters

- *graph-name* (required): The name of the graph.
- *definition-name* (required): The name of the edge collection used in the definition.

Return Codes

- *201*: Returned if the request was successful and waitForSync is true.
- *202*: Returned if the request was successful but waitForSync is false.
- *400*: Returned if no edge definition with this name is found in the graph.
- *404*: Returned if no graph with this name could be found.

Examples

```
shell> curl -X PUT --data-binary @- --dump -
http://localhost:8529/_api/gharial/social/edge/relation <<EOF
{
  "collection" : "relation",
  "from" : [
    "female",
    "male",
    "animal"
  ],
  "to" : [
    "female",
    "male",
    "animal"
  ]
}
EOF

HTTP/1.1 202 Accepted
etag: _U2FBZKu---
content-type: application/json; charset=utf-8
```

show response body

Remove an edge definition from the graph

Remove an edge definition form the graph

`DELETE /_api/gharial/{graph-name}/edge/{definition-name}`

Remove one edge definition from the graph. This will only remove the edge collection, the vertex collections remain untouched and can still be used in your queries.

Example:

```
shell> curl -X DELETE --dump - http://localhost:8529/_api/gharial/social/edge/relation

HTTP/1.1 202 Accepted
etag: _U2FBUQ2---
content-type: application/json; charset=utf-8

{
  "error" : false,
  "graph" : {
    "name" : "social",
    "edgeDefinitions" : [ ],
    "orphanCollections" : [
      "female",
      "male"
    ],
    "isSmart" : false,
    "numberOfShards" : 0,
    "smartGraphAttribute" : "",
    "_id" : "_graphs/social",
    "_rev" : "_U2FBUQ2---"
  },
  "code" : 202
}
```

Path Parameters

- *graph-name* (required): The name of the graph.
- *definition-name* (required): The name of the edge collection used in the definition.

Query Parameters

- *dropCollection* (optional): Drop the collection as well. Collection will only be dropped if it is not used in other graphs.

Return Codes

- *201*: Returned if the edge definition could be removed from the graph and waitForSync is true.
- *202*: Returned if the edge definition could be removed from the graph and waitForSync is false.
- *400*: Returned if no edge definition with this name is found in the graph.
- *404*: Returned if no graph with this name could be found.

Examples

```
shell> curl -X DELETE --dump - http://localhost:8529/_api/gharial/social/edge/relation

HTTP/1.1 202 Accepted
etag: _U2FBUQ2---
content-type: application/json; charset=utf-8
```

show response body

Handling Vertices

Examples will explain the REST API to the graph module on the social graph:

Create a vertex

create a new vertex

`POST /_api/gharial/{graph-name}/vertex/{collection-name}`

Adds a vertex to the given collection.

free style json body

Example:

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/gharial/social/vertex/male <<EOF
{
  "name" : "Francis"
}
EOF

HTTP/1.1 202 Accepted
etag: _U2FBRSe---
content-type: application/json; charset=utf-8

{
  "error" : false,
  "vertex" : {
    "_id" : "male/7808",
    "_key" : "7808",
    "_rev" : "_U2FBRSe---"
  },
  "code" : 202
}
```

Path Parameters

- *graph-name* (required): The name of the graph.
- *collection-name* (required): The name of the vertex collection the vertex belongs to.

Query Parameters

- *waitForSync* (optional): Define if the request should wait until synced to disk.

Post Body

storeThisObject (required): The body has to be the JSON object to be stored.

Return Codes

- *201*: Returned if the vertex could be added and waitForSync is true.
- *202*: Returned if the request was successful but waitForSync is false.
- *404*: Returned if no graph or no vertex collection with this name could be found.

Examples

```
shell> curl -X POST --data-binary @- --dump -
http://localhost:8529/_api/gharial/social/vertex/male <<EOF
{
  "name" : "Francis"
}
EOF

HTTP/1.1 202 Accepted
etag: _U2FBRSe---
content-type: application/json; charset=utf-8
```

show response body

Get a vertex

fetches an existing vertex

`GET /_api/gharial/{graph-name}/vertex/{collection-name}/{vertex-key}`

Gets a vertex from the given collection.

Example:

```
shell> curl --dump - http://localhost:8529/_api/gharial/social/vertex/female/alice

HTTP/1.1 200 OK
etag: _U2FBVSy---
content-type: application/json; charset=utf-8

{
  "error" : false,
  "vertex" : {
    "_key" : "alice",
    "_id" : "female/alice",
    "_rev" : "_U2FBVSy---",
    "name" : "Alice"
  },
  "code" : 200
}
```

Path Parameters

- *graph-name* (required): The name of the graph.
- *collection-name* (required): The name of the vertex collection the vertex belongs to.
- *vertex-key* (required): The *_key* attribute of the vertex.

Header Parameters

- *if-match* (optional): If the "If-Match" header is given, then it must contain exactly one etag. The document is returned, if it has the same revision as the given etag. Otherwise a HTTP 412 is returned. As an alternative you can supply the etag in an attribute rev in the URL.

Return Codes

- *200*: Returned if the vertex could be found.
- *404*: Returned if no graph with this name, no vertex collection or no vertex with this id could be found.
- *412*: Returned if if-match header is given, but the documents revision is different.

Examples

```
shell> curl --dump - http://localhost:8529/_api/gharial/social/vertex/female/alice

HTTP/1.1 200 OK
etag: _U2FBVSy---
content-type: application/json; charset=utf-8
```

show response body

Modify a vertex

replace an existing vertex

`PATCH /_api/gharial/{graph-name}/vertex/{collection-name}/{vertex-key}`

Updates the data of the specific vertex in the collection.

free style json body

Example:

```
shell> curl -X PATCH --data-binary @- --dump - http://localhost:8529/_api/gharial/social/vertex/female/alice <<EOF
{
  "age" : 26
}
EOF

HTTP/1.1 202 Accepted
etag: _U2FBXbW---
content-type: application/json; charset=utf-8

{
  "error" : false,
  "vertex" : {
    "_id" : "female/alice",
    "_key" : "alice",
    "_rev" : "_U2FBXbW---",
    "_oldRev" : "_U2FBXa2---"
  },
  "code" : 202
}
```

Path Parameters

- *graph-name* (required): The name of the graph.
- *collection-name* (required): The name of the vertex collection the vertex belongs to.
- *vertex-key* (required): The *_key* attribute of the vertex.

Query Parameters

- *waitForSync* (optional): Define if the request should wait until synced to disk.
- *keepNull* (optional): Define if values set to null should be stored. By default the key is not removed from the document.

Header Parameters

- *if-match* (optional): If the "If-Match" header is given, then it must contain exactly one etag. The document is updated, if it has the same revision as the given etag. Otherwise a HTTP 412 is returned. As an alternative you can supply the etag in an attribute rev in the URL.

Post Body

replaceAttributes (required): The body has to contain a JSON object containing exactly the attributes that should be replaced.

Return Codes

- *200*: Returned if the vertex could be updated.
- *202*: Returned if the request was successful but waitForSync is false.
- *404*: Returned if no graph with this name, no vertex collection or no vertex with this id could be found.
- *412*: Returned if if-match header is given, but the documents revision is different.

Examples

```
shell> curl -X PATCH --data-binary @- --dump -
http://localhost:8529/_api/gharial/social/vertex/female/alice <<EOF
{
  "age" : 26
}
EOF

HTTP/1.1 202 Accepted
etag: _U2FBXbW---
content-type: application/json; charset=utf-8
```

show response body

Replace a vertex

replaces an existing vertex

`PUT /_api/gharial/{graph-name}/vertex/{collection-name}/{vertex-key}`

Replaces the data of a vertex in the collection.

free style json body

Example:

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/gharial/social/vertex/female/alice <<EOF
{
  "name" : "Alice Cooper",
  "age" : 26
}
EOF

HTTP/1.1 202 Accepted
etag: _U2FBZmu---
content-type: application/json; charset=utf-8

{
  "error" : false,
  "vertex" : {
    "_id" : "female/alice",
    "_key" : "alice",
    "_rev" : "_U2FBZmu---",
    "_oldRev" : "_U2FBZmO---"
  },
  "code" : 202
}
```

Path Parameters

- *graph-name* (required): The name of the graph.
- *collection-name* (required): The name of the vertex collection the vertex belongs to.
- *vertex-key* (required): The *_key* attribute of the vertex.

Query Parameters

- *waitForSync* (optional): Define if the request should wait until synced to disk.

Header Parameters

- *if-match* (optional): If the "If-Match" header is given, then it must contain exactly one etag. The document is updated, if it has the same revision as the given etag. Otherwise a HTTP 412 is returned. As an alternative you can supply the etag in an attribute rev in the URL.

Post Body

storeThisJsonObject (required): The body has to be the JSON object to be stored.

Return Codes

- *200*: Returned if the vertex could be replaced.
- *202*: Returned if the request was successful but waitForSync is false.
- *404*: Returned if no graph with this name, no vertex collection or no vertex with this id could be found.
- *412*: Returned if if-match header is given, but the documents revision is different.

Examples

```
shell> curl -X PUT --data-binary @- --dump -
http://localhost:8529/_api/gharial/social/vertex/female/alice <<EOF
{
  "name" : "Alice Cooper",
  "age" : 26
}
EOF

HTTP/1.1 202 Accepted
etag: _U2FBZmu---
content-type: application/json; charset=utf-8
```

show response body

Remove a vertex

removes a vertex from a graph

`DELETE /_api/gharial/{graph-name}/vertex/{collection-name}/{vertex-key}`

Removes a vertex from the collection.

Example:

```
shell> curl -X DELETE --dump - http://localhost:8529/_api/gharial/social/vertex/female/alice

HTTP/1.1 202 Accepted
content-type: application/json; charset=utf-8

{
  "error" : false,
  "removed" : true,
  "code" : 202
}
```

Path Parameters

- *graph-name* (required): The name of the graph.
- *collection-name* (required): The name of the vertex collection the vertex belongs to.
- *vertex-key* (required): The *_key* attribute of the vertex.

Query Parameters

- *waitForSync* (optional): Define if the request should wait until synced to disk.

Header Parameters

- *if-match* (optional): If the "If-Match" header is given, then it must contain exactly one etag. The document is updated, if it has the same revision as the given etag. Otherwise a HTTP 412 is returned. As an alternative you can supply the etag in an attribute rev in the URL.

Return Codes

- *200*: Returned if the vertex could be removed.
- *202*: Returned if the request was successful but waitForSync is false.
- *404*: Returned if no graph with this name, no vertex collection or no vertex with this id could be found.
- *412*: Returned if if-match header is given, but the documents revision is different.

Examples

```
shell> curl -X DELETE --dump -
http://localhost:8529/_api/gharial/social/vertex/female/alice

HTTP/1.1 202 Accepted
content-type: application/json; charset=utf-8
```

show response body

Handling Edges

Examples will explain the REST API for manipulating edges of the graph module on the knows graph:

Create an edge

Creates an edge in an existing graph

`POST /_api/gharial/{graph-name}/edge/{collection-name}`

Creates a new edge in the collection. Within the body the has to contain a *_from* and *_to* value referencing to valid vertices in the graph. Furthermore the edge has to be valid in the definition of this edge collection.

free style json body

Example:

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/gharial/social/edge/relation <<EOF
{
  "type" : "friend",
  "_from" : "female/alice",
  "_to" : "female/diana"
}
EOF

HTTP/1.1 202 Accepted
etag: _U2FBQP----
content-type: application/json; charset=utf-8

{
  "error" : false,
  "edge" : {
    "_id" : "relation/7696",
    "_key" : "7696",
    "_rev" : "_U2FBQP----"
  },
  "code" : 202
}
```

Path Parameters

- *graph-name* (required): The name of the graph.
- *collection-name* (required): The name of the edge collection the edge belongs to.

Query Parameters

- *waitForSync* (optional): Define if the request should wait until synced to disk.
- *_from* (required):
- *_to* (required):

Post Body

storeThisJsonObject (required): The body has to be the JSON object to be stored.

Return Codes

- *201*: Returned if the edge could be created.
- *202*: Returned if the request was successful but waitForSync is false.
- *404*: Returned if no graph with this name, no edge collection or no edge with this id could be found.

Examples

```
shell> curl -X POST --data-binary @- --dump -
http://localhost:8529/_api/gharial/social/edge/relation <<EOF
{
  "type" : "friend",
  "_from" : "female/alice",
  "_to" : "female/diana"
}
EOF

HTTP/1.1 202 Accepted
etag: _U2FBQP----
content-type: application/json; charset=utf-8
```

show response body

Get an edge

fetch an edge

`GET /_api/gharial/{graph-name}/edge/{collection-name}/{edge-key}`

Gets an edge from the given collection.

Example:

```
shell> curl --dump - http://localhost:8529/_api/gharial/social/edge/relation/aliceAndBob

HTTP/1.1 200 OK
etag: _U2FBUnq---
content-type: application/json; charset=utf-8

{
  "error" : false,
  "edge" : {
    "_key" : "aliceAndBob",
    "_id" : "relation/aliceAndBob",
    "_from" : "female/alice",
    "_to" : "male/bob",
    "_rev" : "_U2FBUnq---",
    "type" : "married"
  },
  "code" : 200
}
```

Path Parameters

- *graph-name* (required): The name of the graph.
- *collection-name* (required): The name of the edge collection the edge belongs to.
- *edge-key* (required): The *_key* attribute of the vertex.

Header Parameters

- *if-match* (optional): If the "If-Match" header is given, then it must contain exactly one etag. The document is returned, if it has the same revision as the given etag. Otherwise a HTTP 412 is returned. As an alternative you can supply the etag in an attribute rev in the URL.

Return Codes

- *200*: Returned if the edge could be found.
- *404*: Returned if no graph with this name, no edge collection or no edge with this id could be found.
- *412*: Returned if if-match header is given, but the documents revision is different.

Examples

```
shell> curl --dump - http://localhost:8529/_api/gharial/social/edge/relation/aliceAndBob

HTTP/1.1 200 OK
etag: _U2FBUnq---
content-type: application/json; charset=utf-8
```

show response body

Examples will explain the API on the social graph:

Modify an edge

modify an existing edge

`PATCH /_api/gharial/{graph-name}/edge/{collection-name}/{edge-key}`

Updates the data of the specific edge in the collection.

free style json body

Example:

```
shell> curl -X PATCH --data-binary @- --dump - http://localhost:8529/_api/gharial/social/edge/relation/aliceAndBob <<
EOF
{
  "since" : "01.01.2001"
}
EOF

HTTP/1.1 202 Accepted
etag: _U2FBXvq---
content-type: application/json; charset=utf-8

{
  "error" : false,
  "edge" : {
    "_id" : "relation/aliceAndBob",
    "_key" : "aliceAndBob",
    "_rev" : "_U2FBXvq---",
    "_oldRev" : "_U2FBXv----"
  },
  "code" : 202
}
```

Path Parameters

- *graph-name* (required): The name of the graph.
- *collection-name* (required): The name of the edge collection the edge belongs to.
- *edge-key* (required): The *_key* attribute of the vertex.

Query Parameters

- *waitForSync* (optional): Define if the request should wait until synced to disk.
- *keepNull* (optional): Define if values set to null should be stored. By default the key is not removed from the document.

Post Body

updateAttributes (required): The body has to be a JSON object containing the attributes to be updated.

Return Codes

- *200*: Returned if the edge could be updated.
- *202*: Returned if the request was successful but waitForSync is false.
- *404*: Returned if no graph with this name, no edge collection or no edge with this id could be found.

Examples

```
shell> curl -X PATCH --data-binary @- --dump -
http://localhost:8529/_api/gharial/social/edge/relation/aliceAndBob <<EOF
{
  "since" : "01.01.2001"
}
EOF

HTTP/1.1 202 Accepted
etag: _U2FBXvq---
content-type: application/json; charset=utf-8
```

show response body

Replace an edge

replace the content of an existing edge

`PUT /_api/gharial/{graph-name}/edge/{collection-name}/{edge-key}`

Replaces the data of an edge in the collection.

free style json body

Example:

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/gharial/social/edge/relation/aliceAndBob <<EO
F
{
  "type" : "divorced",
  "_from" : "female/alice",
  "_to" : "male/bob"
}
EOF

HTTP/1.1 202 Accepted
etag: _U2FBYEa---
content-type: application/json; charset=utf-8

{
  "error" : false,
  "edge" : {
    "_id" : "relation/aliceAndBob",
    "_key" : "aliceAndBob",
    "_rev" : "_U2FBYEa---",
    "_oldRev" : "_U2FBYEC---"
  },
  "code" : 202
}
```

Path Parameters

- *graph-name* (required): The name of the graph.
- *collection-name* (required): The name of the edge collection the edge belongs to.
- *edge-key* (required): The *_key* attribute of the vertex.

Query Parameters

- *waitForSync* (optional): Define if the request should wait until synced to disk.

Header Parameters

- *if-match* (optional): If the "If-Match" header is given, then it must contain exactly one etag. The document is updated, if it has the same revision as the given etag. Otherwise a HTTP 412 is returned. As an alternative you can supply the etag in an attribute rev in the URL.

Post Body

storeThisJsonObject (required): The body has to be the JSON object to be stored.

Return Codes

- *201*: Returned if the request was successful but waitForSync is true.
- *202*: Returned if the request was successful but waitForSync is false.
- *404*: Returned if no graph with this name, no edge collection or no edge with this id could be found.
- *412*: Returned if if-match header is given, but the documents revision is different.

Examples

```
shell> curl -X PUT --data-binary @- --dump -
http://localhost:8529/_api/gharial/social/edge/relation/aliceAndBob <<EOF
{
  "type" : "divorced",
  "_from" : "female/alice",
  "_to" : "male/bob"
}
EOF

HTTP/1.1 202 Accepted
etag: _U2FBYEa---
content-type: application/json; charset=utf-8
```

show response body

Remove an edge

removes an edge from graph

`DELETE /_api/gharial/{graph-name}/edge/{collection-name}/{edge-key}`

Removes an edge from the collection.

Example:

```
shell> curl -X DELETE --dump - http://localhost:8529/_api/gharial/social/edge/relation/aliceAndBob

HTTP/1.1 202 Accepted
content-type: application/json; charset=utf-8

{
  "error" : false,
  "removed" : true,
  "code" : 202
}
```

Path Parameters

- *graph-name* (required): The name of the graph.
- *collection-name* (required): The name of the edge collection the edge belongs to.

- *edge-key* (required): The *_key* attribute of the vertex.

Query Parameters

- *waitForSync* (optional): Define if the request should wait until synced to disk.

Header Parameters

- *if-match* (optional): If the "If-Match" header is given, then it must contain exactly one etag. The document is updated, if it has the same revision as the given etag. Otherwise a HTTP 412 is returned. As an alternative you can supply the etag in an attribute rev in the URL.

Return Codes

- *200*: Returned if the edge could be removed.
- *202*: Returned if the request was successful but waitForSync is false.
- *404*: Returned if no graph with this name, no edge collection or no edge with this id could be found.
- *412*: Returned if if-match header is given, but the documents revision is different.

Examples

```
shell> curl -X DELETE --dump -
http://localhost:8529/_api/gharial/social/edge/relation/aliceAndBob

HTTP/1.1 202 Accepted
content-type: application/json; charset=utf-8
```

show response body

HTTP Interface for Traversals

Traversals

ArangoDB's graph traversals are executed on the server. Traversals can be initiated by clients by sending the traversal description for execution to the server.

Traversals in ArangoDB are used to walk over a graph stored in one edge collection. It can easily be described which edges of the graph should be followed and which actions should be performed on each visited vertex. Furthermore the ordering of visiting the nodes can be specified, for instance depth-first or breadth-first search are offered.

Executing Traversals via HTTP

executes a traversal

execute a server-side traversal

`POST /_api/traversal`

Starts a traversal starting from a given vertex and following. edges contained in a given edgeCollection. The request must contain the following attributes.

A json post document with these Properties is required:

- **sort**: body (JavaScript) code of a custom comparison function for the edges. The signature of this function is *(l, r) -> integer* (where l and r are edges) and must return -1 if l is smaller than, +1 if l is greater than, and 0 if l and r are equal. The reason for this is the following: The order of edges returned for a certain vertex is undefined. This is because there is no natural order of edges for a vertex with multiple connected edges. To explicitly define the order in which edges on the vertex are followed, you can specify an edge comparator function with this attribute. Note that the value here has to be a string to conform to the JSON standard, which in turn is parsed as function body on the server side. Furthermore note that this attribute is only used for the standard expanders. If you use your custom expander you have to do the sorting yourself within the expander code.
- **direction**: direction for traversal
  - *if set*, must be either *"outbound"*, *"inbound"*, or *"any"*
  - *if not set*, the *expander* attribute must be specified
- **minDepth**: ANDed with any existing filters): visits only nodes in at least the given depth
- **startVertex**: id of the startVertex, e.g. *"users/foo"*.
- **visitor**: body (JavaScript) code of custom visitor function function signature: *(config, result, vertex, path, connected) -> void* The visitor function can do anything, but its return value is ignored. To populate a result, use the *result* variable by reference. Note that the *connected* argument is only populated when the *order* attribute is set to *"preorder-expander"*.
- **itemOrder**: item iteration order can be *"forward"* or *"backward"*
- **strategy**: traversal strategy can be *"depthfirst"* or *"breadthfirst"*
- **filter**: default is to include all nodes: body (JavaScript code) of custom filter function function signature: *(config, vertex, path) -> mixed* can return four different string values:
  - *"exclude"* -> this vertex will not be visited.
  - *"prune"* -> the edges of this vertex will not be followed.
  - *""* or *undefined* -> visit the vertex and follow it's edges.
  - *Array* -> containing any combination of the above. If there is at least one *"exclude"* or *"prune"* respectivly is contained, it's effect will occur.
- **init**: body (JavaScript) code of custom result initialization function function signature: *(config, result) -> void* initialize any values in result with what is required
- **maxIterations**: Maximum number of iterations in each traversal. This number can be set to prevent endless loops in traversal of cyclic graphs. When a traversal performs as many iterations as the *maxIterations* value, the traversal will abort with an error. If *maxIterations* is not set, a server-defined value may be used.
- **maxDepth**: ANDed with any existing filters visits only nodes in at most the given depth
- **uniqueness**: specifies uniqueness for vertices and edges visited. If set, must be an object like this: `"uniqueness": {"vertices": "none"|"global"|"path", "edges": "none"|"global"|"path"}`

- **order**: traversal order can be *"preorder"*, *"postorder"* or *"preorder-expander"*
- **graphName**: name of the graph that contains the edges. Either *edgeCollection* or *graphName* has to be given. In case both values are set the *graphName* is prefered.
- **expander**: body (JavaScript) code of custom expander function *must* be set if *direction* attribute is **not** set function signature: *(config, vertex, path) -> array* expander must return an array of the connections for *vertex* each connection is an object with the attributes *edge* and *vertex*
- **edgeCollection**: name of the collection that contains the edges.

If the Traversal is successfully executed *HTTP 200* will be returned. Additionally the *result* object will be returned by the traversal.

For successful traversals, the returned JSON object has the following properties:

- *error*: boolean flag to indicate if an error occurred (*false* in this case)
- *code*: the HTTP status code
- *result*: the return value of the traversal

If the traversal specification is either missing or malformed, the server will respond with *HTTP 400*.

The body of the response will then contain a JSON object with additional error details. The object has the following attributes:

- *error*: boolean flag to indicate that an error occurred (*true* in this case)
- *code*: the HTTP status code
- *errorNum*: the server error number
- *errorMessage*: a descriptive error message

Example:

In the following examples the underlying graph will contain five persons *Alice*, *Bob*, *Charlie*, *Dave* and *Eve*. We will have the following directed relations:

- *Alice* knows *Bob*
- *Bob* knows *Charlie*
- *Bob* knows *Dave*
- *Eve* knows *Alice*
- *Eve* knows *Bob*

The starting vertex will always be Alice.

Follow only outbound edges

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/traversal <<EOF
{
  "startVertex" : "persons/alice",
  "graphName" : "knows_graph",
  "direction" : "outbound"
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "result" : {
    "visited" : {
      "vertices" : [
        {
          "_key" : "alice",
          "_id" : "persons/alice",
          "_rev" : "_U2FCvNm---",
          "name" : "Alice"
        },
        {
          "_key" : "bob",
```

```
          "_id" : "persons/bob",
          "_rev" : "_U2FCvNm--_",
          "name" : "Bob"
        },
        {
          "_key" : "charlie",
          "_id" : "persons/charlie",
          "_rev" : "_U2FCvNq---",
          "name" : "Charlie"
        },
        {
          "_key" : "dave",
          "_id" : "persons/dave",
          "_rev" : "_U2FCvNq--_",
          "name" : "Dave"
        }
      ],
      "paths" : [
        {
          "edges" : [ ],
          "vertices" : [
            {
              "_key" : "alice",
              "_id" : "persons/alice",
              "_rev" : "_U2FCvNm---",
              "name" : "Alice"
            }
          ]
        },
        {
          "edges" : [
            {
              "_key" : "12019",
              "_id" : "knows/12019",
              "_from" : "persons/alice",
              "_to" : "persons/bob",
              "_rev" : "_U2FCvNu---"
            }
          ],
          "vertices" : [
            {
              "_key" : "alice",
              "_id" : "persons/alice",
              "_rev" : "_U2FCvNm---",
              "name" : "Alice"
            },
            {
              "_key" : "bob",
              "_id" : "persons/bob",
              "_rev" : "_U2FCvNm--_",
              "name" : "Bob"
            }
          ]
        },
        {
          "edges" : [
            {
              "_key" : "12019",
              "_id" : "knows/12019",
              "_from" : "persons/alice",
              "_to" : "persons/bob",
              "_rev" : "_U2FCvNu---"
            },
            {
              "_key" : "12023",
              "_id" : "knows/12023",
              "_from" : "persons/bob",
              "_to" : "persons/charlie",
              "_rev" : "_U2FCvNu--_"
            }
          ],
          "vertices" : [
            {
              "_key" : "alice",
              "_id" : "persons/alice",
              "_rev" : "_U2FCvNm---",
```

```
              "name" : "Alice"
            },
            {
              "_key" : "bob",
              "_id" : "persons/bob",
              "_rev" : "_U2FCvNm--_",
              "name" : "Bob"
            },
            {
              "_key" : "charlie",
              "_id" : "persons/charlie",
              "_rev" : "_U2FCvNq---",
              "name" : "Charlie"
            }
          ]
        },
        {
          "edges" : [
            {
              "_key" : "12019",
              "_id" : "knows/12019",
              "_from" : "persons/alice",
              "_to" : "persons/bob",
              "_rev" : "_U2FCvNu---"
            },
            {
              "_key" : "12026",
              "_id" : "knows/12026",
              "_from" : "persons/bob",
              "_to" : "persons/dave",
              "_rev" : "_U2FCvNu--A"
            }
          ],
          "vertices" : [
            {
              "_key" : "alice",
              "_id" : "persons/alice",
              "_rev" : "_U2FCvNm---",
              "name" : "Alice"
            },
            {
              "_key" : "bob",
              "_id" : "persons/bob",
              "_rev" : "_U2FCvNm--_",
              "name" : "Bob"
            },
            {
              "_key" : "dave",
              "_id" : "persons/dave",
              "_rev" : "_U2FCvNq--_",
              "name" : "Dave"
            }
          ]
        }
      ]
    }
  },
  "error" : false,
  "code" : 200
}
```

Example:

Follow only inbound edges

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/traversal <<EOF
{
  "startVertex" : "persons/alice",
  "graphName" : "knows_graph",
  "direction" : "inbound"
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "result" : {
    "visited" : {
      "vertices" : [
        {
          "_key" : "alice",
          "_id" : "persons/alice",
          "_rev" : "_U2FCtze---",
          "name" : "Alice"
        },
        {
          "_key" : "eve",
          "_id" : "persons/eve",
          "_rev" : "_U2FCtzi--_",
          "name" : "Eve"
        }
      ],
      "paths" : [
        {
          "edges" : [ ],
          "vertices" : [
            {
              "_key" : "alice",
              "_id" : "persons/alice",
              "_rev" : "_U2FCtze---",
              "name" : "Alice"
            }
          ]
        },
        {
          "edges" : [
            {
              "_key" : "11774",
              "_id" : "knows/11774",
              "_from" : "persons/eve",
              "_to" : "persons/alice",
              "_rev" : "_U2FCtzq--_"
            }
          ],
          "vertices" : [
            {
              "_key" : "alice",
              "_id" : "persons/alice",
              "_rev" : "_U2FCtze---",
              "name" : "Alice"
            },
            {
              "_key" : "eve",
              "_id" : "persons/eve",
              "_rev" : "_U2FCtzi--_",
              "name" : "Eve"
            }
          ]
        }
      ]
    }
  },
  "error" : false,
  "code" : 200
}
```

Example:

Follow any direction of edges

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/traversal <<EOF
{
  "startVertex" : "persons/alice",
  "graphName" : "knows_graph",
  "direction" : "any",
  "uniqueness" : {
    "vertices" : "none",
    "edges" : "global"
  }
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "result" : {
    "visited" : {
      "vertices" : [
        {
          "_key" : "alice",
          "_id" : "persons/alice",
          "_rev" : "_U2FCsUK---",
          "name" : "Alice"
        },
        {
          "_key" : "bob",
          "_id" : "persons/bob",
          "_rev" : "_U2FCsUO---",
          "name" : "Bob"
        },
        {
          "_key" : "charlie",
          "_id" : "persons/charlie",
          "_rev" : "_U2FCsUO--_",
          "name" : "Charlie"
        },
        {
          "_key" : "dave",
          "_id" : "persons/dave",
          "_rev" : "_U2FCsUS---",
          "name" : "Dave"
        },
        {
          "_key" : "eve",
          "_id" : "persons/eve",
          "_rev" : "_U2FCsUS--_",
          "name" : "Eve"
        },
        {
          "_key" : "alice",
          "_id" : "persons/alice",
          "_rev" : "_U2FCsUK---",
          "name" : "Alice"
        }
      ],
      "paths" : [
        {
          "edges" : [ ],
          "vertices" : [
            {
              "_key" : "alice",
              "_id" : "persons/alice",
              "_rev" : "_U2FCsUK---",
              "name" : "Alice"
            }
          ]
        },
        {
          "edges" : [
            {
              "_key" : "11341",
              "_id" : "knows/11341",
```

```
            "_from" : "persons/alice",
            "_to" : "persons/bob",
            "_rev" : "_U2FCsUW---"
          }
        ],
        "vertices" : [
          {
            "_key" : "alice",
            "_id" : "persons/alice",
            "_rev" : "_U2FCsUK---",
            "name" : "Alice"
          },
          {
            "_key" : "bob",
            "_id" : "persons/bob",
            "_rev" : "_U2FCsUO---",
            "name" : "Bob"
          }
        ]
      },
      {
        "edges" : [
          {
            "_key" : "11341",
            "_id" : "knows/11341",
            "_from" : "persons/alice",
            "_to" : "persons/bob",
            "_rev" : "_U2FCsUW---"
          },
          {
            "_key" : "11345",
            "_id" : "knows/11345",
            "_from" : "persons/bob",
            "_to" : "persons/charlie",
            "_rev" : "_U2FCsUa---"
          }
        ],
        "vertices" : [
          {
            "_key" : "alice",
            "_id" : "persons/alice",
            "_rev" : "_U2FCsUK---",
            "name" : "Alice"
          },
          {
            "_key" : "bob",
            "_id" : "persons/bob",
            "_rev" : "_U2FCsUO---",
            "name" : "Bob"
          },
          {
            "_key" : "charlie",
            "_id" : "persons/charlie",
            "_rev" : "_U2FCsUO--_",
            "name" : "Charlie"
          }
        ]
      },
      {
        "edges" : [
          {
            "_key" : "11341",
            "_id" : "knows/11341",
            "_from" : "persons/alice",
            "_to" : "persons/bob",
            "_rev" : "_U2FCsUW---"
          },
          {
            "_key" : "11348",
            "_id" : "knows/11348",
            "_from" : "persons/bob",
            "_to" : "persons/dave",
            "_rev" : "_U2FCsUa--_"
          }
        ],
        "vertices" : [
```

```
            {
              "_key" : "alice",
              "_id" : "persons/alice",
              "_rev" : "_U2FCsUK---",
              "name" : "Alice"
            },
            {
              "_key" : "bob",
              "_id" : "persons/bob",
              "_rev" : "_U2FCsUO---",
              "name" : "Bob"
            },
            {
              "_key" : "dave",
              "_id" : "persons/dave",
              "_rev" : "_U2FCsUS---",
              "name" : "Dave"
            }
          ]
        },
        {
          "edges" : [
            {
              "_key" : "11341",
              "_id" : "knows/11341",
              "_from" : "persons/alice",
              "_to" : "persons/bob",
              "_rev" : "_U2FCsUW---"
            },
            {
              "_key" : "11354",
              "_id" : "knows/11354",
              "_from" : "persons/eve",
              "_to" : "persons/bob",
              "_rev" : "_U2FCsUe--_"
            }
          ],
          "vertices" : [
            {
              "_key" : "alice",
              "_id" : "persons/alice",
              "_rev" : "_U2FCsUK---",
              "name" : "Alice"
            },
            {
              "_key" : "bob",
              "_id" : "persons/bob",
              "_rev" : "_U2FCsUO---",
              "name" : "Bob"
            },
            {
              "_key" : "eve",
              "_id" : "persons/eve",
              "_rev" : "_U2FCsUS--_",
              "name" : "Eve"
            }
          ]
        },
        {
          "edges" : [
            {
              "_key" : "11341",
              "_id" : "knows/11341",
              "_from" : "persons/alice",
              "_to" : "persons/bob",
              "_rev" : "_U2FCsUW---"
            },
            {
              "_key" : "11354",
              "_id" : "knows/11354",
              "_from" : "persons/eve",
              "_to" : "persons/bob",
              "_rev" : "_U2FCsUe--_"
            },
            {
              "_key" : "11351",
```

```
              "_id" : "knows/11351",
              "_from" : "persons/eve",
              "_to" : "persons/alice",
              "_rev" : "_U2FCsUe---"
            }
          ],
          "vertices" : [
            {
              "_key" : "alice",
              "_id" : "persons/alice",
              "_rev" : "_U2FCsUK---",
              "name" : "Alice"
            },
            {
              "_key" : "bob",
              "_id" : "persons/bob",
              "_rev" : "_U2FCsUO---",
              "name" : "Bob"
            },
            {
              "_key" : "eve",
              "_id" : "persons/eve",
              "_rev" : "_U2FCsUS--_",
              "name" : "Eve"
            },
            {
              "_key" : "alice",
              "_id" : "persons/alice",
              "_rev" : "_U2FCsUK---",
              "name" : "Alice"
            }
          ]
        }
      ]
    }
  },
  "error" : false,
  "code" : 200
}
```

Example:

Excluding *Charlie* and *Bob*

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/traversal <<EOF
{
  "startVertex" : "persons/alice",
  "graphName" : "knows_graph",
  "direction" : "outbound",
  "filter" : "if (vertex.name === \"Bob\" ||     vertex.name === \"Charlie\") {  return \"exclude\";}return;"
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "result" : {
    "visited" : {
      "vertices" : [
        {
          "_key" : "alice",
          "_id" : "persons/alice",
          "_rev" : "_U2FCtV----",
          "name" : "Alice"
        },
        {
          "_key" : "dave",
          "_id" : "persons/dave",
          "_rev" : "_U2FCtVC--_",
          "name" : "Dave"
        }
```

```
      ],
      "paths" : [
        {
          "edges" : [ ],
          "vertices" : [
            {
              "_key" : "alice",
              "_id" : "persons/alice",
              "_rev" : "_U2FCtV----",
              "name" : "Alice"
            }
          ]
        },
        {
          "edges" : [
            {
              "_key" : "11651",
              "_id" : "knows/11651",
              "_from" : "persons/alice",
              "_to" : "persons/bob",
              "_rev" : "_U2FCtVG--_"
            },
            {
              "_key" : "11658",
              "_id" : "knows/11658",
              "_from" : "persons/bob",
              "_to" : "persons/dave",
              "_rev" : "_U2FCtVK---"
            }
          ],
          "vertices" : [
            {
              "_key" : "alice",
              "_id" : "persons/alice",
              "_rev" : "_U2FCtV----",
              "name" : "Alice"
            },
            {
              "_key" : "bob",
              "_id" : "persons/bob",
              "_rev" : "_U2FCtV---_",
              "name" : "Bob"
            },
            {
              "_key" : "dave",
              "_id" : "persons/dave",
              "_rev" : "_U2FCtVC--_",
              "name" : "Dave"
            }
          ]
        }
      ]
    }
  },
  "error" : false,
  "code" : 200
}
```


Example:

Do not follow edges from *Bob*

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/traversal <<EOF
{
  "startVertex" : "persons/alice",
  "graphName" : "knows_graph",
  "direction" : "outbound",
  "filter" : "if (vertex.name === \"Bob\") {return \"prune\";}return;"
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "result" : {
    "visited" : {
      "vertices" : [
        {
          "_key" : "alice",
          "_id" : "persons/alice",
          "_rev" : "_U2FCtkS---",
          "name" : "Alice"
        },
        {
          "_key" : "bob",
          "_id" : "persons/bob",
          "_rev" : "_U2FCtkS--_",
          "name" : "Bob"
        }
      ],
      "paths" : [
        {
          "edges" : [ ],
          "vertices" : [
            {
              "_key" : "alice",
              "_id" : "persons/alice",
              "_rev" : "_U2FCtkS---",
              "name" : "Alice"
            }
          ]
        },
        {
          "edges" : [
            {
              "_key" : "11710",
              "_id" : "knows/11710",
              "_from" : "persons/alice",
              "_to" : "persons/bob",
              "_rev" : "_U2FCtka---"
            }
          ],
          "vertices" : [
            {
              "_key" : "alice",
              "_id" : "persons/alice",
              "_rev" : "_U2FCtkS---",
              "name" : "Alice"
            },
            {
              "_key" : "bob",
              "_id" : "persons/bob",
              "_rev" : "_U2FCtkS--_",
              "name" : "Bob"
            }
          ]
        }
      ]
    }
  },
  "error" : false,
  "code" : 200
}
```

Example:

Visit only nodes in a depth of at least 2

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/traversal <<EOF
{
  "startVertex" : "persons/alice",
  "graphName" : "knows_graph",
  "direction" : "outbound",
  "minDepth" : 2
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "result" : {
    "visited" : {
      "vertices" : [
        {
          "_key" : "charlie",
          "_id" : "persons/charlie",
          "_rev" : "_U2FCun----",
          "name" : "Charlie"
        },
        {
          "_key" : "dave",
          "_id" : "persons/dave",
          "_rev" : "_U2FCun---_",
          "name" : "Dave"
        }
      ],
      "paths" : [
        {
          "edges" : [
            {
              "_key" : "11960",
              "_id" : "knows/11960",
              "_from" : "persons/alice",
              "_to" : "persons/bob",
              "_rev" : "_U2FCunC--_"
            },
            {
              "_key" : "11964",
              "_id" : "knows/11964",
              "_from" : "persons/bob",
              "_to" : "persons/charlie",
              "_rev" : "_U2FCunG---"
            }
          ],
          "vertices" : [
            {
              "_key" : "alice",
              "_id" : "persons/alice",
              "_rev" : "_U2FCum6---",
              "name" : "Alice"
            },
            {
              "_key" : "bob",
              "_id" : "persons/bob",
              "_rev" : "_U2FCum6--_",
              "name" : "Bob"
            },
            {
              "_key" : "charlie",
              "_id" : "persons/charlie",
              "_rev" : "_U2FCun----",
              "name" : "Charlie"
            }
          ]
        },
        {
          "edges" : [
```

```
            {
              "_key" : "11960",
              "_id" : "knows/11960",
              "_from" : "persons/alice",
              "_to" : "persons/bob",
              "_rev" : "_U2FCunC--_"
            },
            {
              "_key" : "11967",
              "_id" : "knows/11967",
              "_from" : "persons/bob",
              "_to" : "persons/dave",
              "_rev" : "_U2FCunG--_"
            }
          ],
          "vertices" : [
            {
              "_key" : "alice",
              "_id" : "persons/alice",
              "_rev" : "_U2FCum6---",
              "name" : "Alice"
            },
            {
              "_key" : "bob",
              "_id" : "persons/bob",
              "_rev" : "_U2FCum6--_",
              "name" : "Bob"
            },
            {
              "_key" : "dave",
              "_id" : "persons/dave",
              "_rev" : "_U2FCun---_",
              "name" : "Dave"
            }
          ]
        }
      ]
    }
  },
  "error" : false,
  "code" : 200
}
```


Example:

Visit only nodes in a depth of at most 1

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/traversal <<EOF
{
  "startVertex" : "persons/alice",
  "graphName" : "knows_graph",
  "direction" : "outbound",
  "maxDepth" : 1
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "result" : {
    "visited" : {
      "vertices" : [
        {
          "_key" : "alice",
          "_id" : "persons/alice",
          "_rev" : "_U2FCuDK---",
          "name" : "Alice"
        },
        {
          "_key" : "bob",
          "_id" : "persons/bob",
          "_rev" : "_U2FCuDO---",
          "name" : "Bob"
        }
      ],
      "paths" : [
        {
          "edges" : [ ],
          "vertices" : [
            {
              "_key" : "alice",
              "_id" : "persons/alice",
              "_rev" : "_U2FCuDK---",
              "name" : "Alice"
            }
          ]
        },
        {
          "edges" : [
            {
              "_key" : "11819",
              "_id" : "knows/11819",
              "_from" : "persons/alice",
              "_to" : "persons/bob",
              "_rev" : "_U2FCuDW---"
            }
          ],
          "vertices" : [
            {
              "_key" : "alice",
              "_id" : "persons/alice",
              "_rev" : "_U2FCuDK---",
              "name" : "Alice"
            },
            {
              "_key" : "bob",
              "_id" : "persons/bob",
              "_rev" : "_U2FCuDO---",
              "name" : "Bob"
            }
          ]
        }
      ]
    }
  },
  "error" : false,
  "code" : 200
}
```

Example:

Using a visitor function to return vertex ids only

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/traversal <<EOF
{
  "startVertex" : "persons/alice",
  "graphName" : "knows_graph",
  "direction" : "outbound",
  "visitor" : "result.visited.vertices.push(vertex._id);"
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "result" : {
    "visited" : {
      "vertices" : [
        "persons/alice",
        "persons/bob",
        "persons/charlie",
        "persons/dave"
      ],
      "paths" : [ ]
    }
  },
  "error" : false,
  "code" : 200
}
```

Example:

Count all visited nodes and return a list of nodes only

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/traversal <<EOF
{
  "startVertex" : "persons/alice",
  "graphName" : "knows_graph",
  "direction" : "outbound",
  "init" : "result.visited = 0; result.myVertices = [ ];",
  "visitor" : "result.visited++; result.myVertices.push(vertex);"
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "result" : {
    "visited" : 4,
    "myVertices" : [
      {
        "_key" : "alice",
        "_id" : "persons/alice",
        "_rev" : "_U2FCvuq---",
        "name" : "Alice"
      },
      {
        "_key" : "bob",
        "_id" : "persons/bob",
        "_rev" : "_U2FCvuu---",
        "name" : "Bob"
      },
      {
        "_key" : "charlie",
        "_id" : "persons/charlie",
        "_rev" : "_U2FCvuu--_",
        "name" : "Charlie"
      },
      {
        "_key" : "dave",
        "_id" : "persons/dave",
        "_rev" : "_U2FCvuy---",
        "name" : "Dave"
      }
    ]
  },
  "error" : false,
  "code" : 200
}
```

Example:

Expand only inbound edges of *Alice* and outbound edges of *Eve*

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/traversal <<EOF
{
  "startVertex" : "persons/alice",
  "graphName" : "knows_graph",
  "expander" : "var connections = [ ];if (vertex.name === \"Alice\") {config.datasource.getInEdges(vertex).forEach(fu
nction (e) {connections.push({ vertex: require(\"internal\").db._document(e._from), edge: e});});}if (vertex.name ===
 \"Eve\") {config.datasource.getOutEdges(vertex).forEach(function (e) {connections.push({vertex: require(\"internal\"
).db._document(e._to), edge: e});});}return connections;"
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "result" : {
    "visited" : {
      "vertices" : [
        {
          "_key" : "alice",
```

```
      "_id" : "persons/alice",
      "_rev" : "_U2FCv82---",
      "name" : "Alice"
    },
    {
      "_key" : "eve",
      "_id" : "persons/eve",
      "_rev" : "_U2FCv9----",
      "name" : "Eve"
    },
    {
      "_key" : "bob",
      "_id" : "persons/bob",
      "_rev" : "_U2FCv82--_",
      "name" : "Bob"
    }
  ],
  "paths" : [
    {
      "edges" : [ ],
      "vertices" : [
        {
          "_key" : "alice",
          "_id" : "persons/alice",
          "_rev" : "_U2FCv82---",
          "name" : "Alice"
        }
      ]
    },
    {
      "edges" : [
        {
          "_key" : "12232",
          "_id" : "knows/12232",
          "_from" : "persons/eve",
          "_to" : "persons/alice",
          "_rev" : "_U2FCv9C--_"
        }
      ],
      "vertices" : [
        {
          "_key" : "alice",
          "_id" : "persons/alice",
          "_rev" : "_U2FCv82---",
          "name" : "Alice"
        },
        {
          "_key" : "eve",
          "_id" : "persons/eve",
          "_rev" : "_U2FCv9----",
          "name" : "Eve"
        }
      ]
    },
    {
      "edges" : [
        {
          "_key" : "12232",
          "_id" : "knows/12232",
          "_from" : "persons/eve",
          "_to" : "persons/alice",
          "_rev" : "_U2FCv9C--_"
        },
        {
          "_key" : "12235",
          "_id" : "knows/12235",
          "_from" : "persons/eve",
          "_to" : "persons/bob",
          "_rev" : "_U2FCv9G---"
        }
      ],
      "vertices" : [
        {
          "_key" : "alice",
          "_id" : "persons/alice",
          "_rev" : "_U2FCv82---",
```

```
              "name" : "Alice"
            },
            {
              "_key" : "eve",
              "_id" : "persons/eve",
              "_rev" : "_U2FCv9----",
              "name" : "Eve"
            },
            {
              "_key" : "bob",
              "_id" : "persons/bob",
              "_rev" : "_U2FCv82--_",
              "name" : "Bob"
            }
          ]
        }
      ]
    }
  },
  "error" : false,
  "code" : 200
}
```

Example:

Follow the *depthfirst* strategy

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/traversal <<EOF
{
  "startVertex" : "persons/alice",
  "graphName" : "knows_graph",
  "direction" : "any",
  "strategy" : "depthfirst"
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "result" : {
    "visited" : {
      "vertices" : [
        {
          "_key" : "alice",
          "_id" : "persons/alice",
          "_rev" : "_U2FCs16--_",
          "name" : "Alice"
        },
        {
          "_key" : "bob",
          "_id" : "persons/bob",
          "_rev" : "_U2FCs2----",
          "name" : "Bob"
        },
        {
          "_key" : "charlie",
          "_id" : "persons/charlie",
          "_rev" : "_U2FCs2---_",
          "name" : "Charlie"
        },
        {
          "_key" : "dave",
          "_id" : "persons/dave",
          "_rev" : "_U2FCs2---A",
          "name" : "Dave"
        },
        {
          "_key" : "eve",
          "_id" : "persons/eve",
          "_rev" : "_U2FCs2---B",
          "name" : "Eve"
```

```
      },
      {
        "_key" : "alice",
        "_id" : "persons/alice",
        "_rev" : "_U2FCs16--_",
        "name" : "Alice"
      },
      {
        "_key" : "eve",
        "_id" : "persons/eve",
        "_rev" : "_U2FCs2---B",
        "name" : "Eve"
      },
      {
        "_key" : "bob",
        "_id" : "persons/bob",
        "_rev" : "_U2FCs2----",
        "name" : "Bob"
      },
      {
        "_key" : "charlie",
        "_id" : "persons/charlie",
        "_rev" : "_U2FCs2---_",
        "name" : "Charlie"
      },
      {
        "_key" : "dave",
        "_id" : "persons/dave",
        "_rev" : "_U2FCs2---A",
        "name" : "Dave"
      },
      {
        "_key" : "alice",
        "_id" : "persons/alice",
        "_rev" : "_U2FCs16--_",
        "name" : "Alice"
      }
    ],
    "paths" : [
      {
        "edges" : [ ],
        "vertices" : [
          {
            "_key" : "alice",
            "_id" : "persons/alice",
            "_rev" : "_U2FCs16--_",
            "name" : "Alice"
          }
        ]
      },
      {
        "edges" : [
          {
            "_key" : "11496",
            "_id" : "knows/11496",
            "_from" : "persons/alice",
            "_to" : "persons/bob",
            "_rev" : "_U2FCs2C---"
          }
        ],
        "vertices" : [
          {
            "_key" : "alice",
            "_id" : "persons/alice",
            "_rev" : "_U2FCs16--_",
            "name" : "Alice"
          },
          {
            "_key" : "bob",
            "_id" : "persons/bob",
            "_rev" : "_U2FCs2----",
            "name" : "Bob"
          }
        ]
      },
      {
```

```
          "edges" : [
            {
              "_key" : "11496",
              "_id" : "knows/11496",
              "_from" : "persons/alice",
              "_to" : "persons/bob",
              "_rev" : "_U2FCs2C---"
            },
            {
              "_key" : "11500",
              "_id" : "knows/11500",
              "_from" : "persons/bob",
              "_to" : "persons/charlie",
              "_rev" : "_U2FCs2C--_"
            }
          ],
          "vertices" : [
            {
              "_key" : "alice",
              "_id" : "persons/alice",
              "_rev" : "_U2FCs16--_",
              "name" : "Alice"
            },
            {
              "_key" : "bob",
              "_id" : "persons/bob",
              "_rev" : "_U2FCs2----",
              "name" : "Bob"
            },
            {
              "_key" : "charlie",
              "_id" : "persons/charlie",
              "_rev" : "_U2FCs2---_",
              "name" : "Charlie"
            }
          ]
        },
        {
          "edges" : [
            {
              "_key" : "11496",
              "_id" : "knows/11496",
              "_from" : "persons/alice",
              "_to" : "persons/bob",
              "_rev" : "_U2FCs2C---"
            },
            {
              "_key" : "11503",
              "_id" : "knows/11503",
              "_from" : "persons/bob",
              "_to" : "persons/dave",
              "_rev" : "_U2FCs2C--A"
            }
          ],
          "vertices" : [
            {
              "_key" : "alice",
              "_id" : "persons/alice",
              "_rev" : "_U2FCs16--_",
              "name" : "Alice"
            },
            {
              "_key" : "bob",
              "_id" : "persons/bob",
              "_rev" : "_U2FCs2----",
              "name" : "Bob"
            },
            {
              "_key" : "dave",
              "_id" : "persons/dave",
              "_rev" : "_U2FCs2---A",
              "name" : "Dave"
            }
          ]
        },
        {
```

```
          "edges" : [
            {
              "_key" : "11496",
              "_id" : "knows/11496",
              "_from" : "persons/alice",
              "_to" : "persons/bob",
              "_rev" : "_U2FCs2C---"
            },
            {
              "_key" : "11509",
              "_id" : "knows/11509",
              "_from" : "persons/eve",
              "_to" : "persons/bob",
              "_rev" : "_U2FCs2C--C"
            }
          ],
          "vertices" : [
            {
              "_key" : "alice",
              "_id" : "persons/alice",
              "_rev" : "_U2FCs16--_",
              "name" : "Alice"
            },
            {
              "_key" : "bob",
              "_id" : "persons/bob",
              "_rev" : "_U2FCs2----",
              "name" : "Bob"
            },
            {
              "_key" : "eve",
              "_id" : "persons/eve",
              "_rev" : "_U2FCs2---B",
              "name" : "Eve"
            }
          ]
        },
        {
          "edges" : [
            {
              "_key" : "11496",
              "_id" : "knows/11496",
              "_from" : "persons/alice",
              "_to" : "persons/bob",
              "_rev" : "_U2FCs2C---"
            },
            {
              "_key" : "11509",
              "_id" : "knows/11509",
              "_from" : "persons/eve",
              "_to" : "persons/bob",
              "_rev" : "_U2FCs2C--C"
            },
            {
              "_key" : "11506",
              "_id" : "knows/11506",
              "_from" : "persons/eve",
              "_to" : "persons/alice",
              "_rev" : "_U2FCs2C--B"
            }
          ],
          "vertices" : [
            {
              "_key" : "alice",
              "_id" : "persons/alice",
              "_rev" : "_U2FCs16--_",
              "name" : "Alice"
            },
            {
              "_key" : "bob",
              "_id" : "persons/bob",
              "_rev" : "_U2FCs2----",
              "name" : "Bob"
            },
            {
              "_key" : "eve",
```

```
            "_id" : "persons/eve",
            "_rev" : "_U2FCs2---B",
            "name" : "Eve"
          },
          {
            "_key" : "alice",
            "_id" : "persons/alice",
            "_rev" : "_U2FCs16--_",
            "name" : "Alice"
          }
        ]
      },
      {
        "edges" : [
          {
            "_key" : "11506",
            "_id" : "knows/11506",
            "_from" : "persons/eve",
            "_to" : "persons/alice",
            "_rev" : "_U2FCs2C--B"
          }
        ],
        "vertices" : [
          {
            "_key" : "alice",
            "_id" : "persons/alice",
            "_rev" : "_U2FCs16--_",
            "name" : "Alice"
          },
          {
            "_key" : "eve",
            "_id" : "persons/eve",
            "_rev" : "_U2FCs2---B",
            "name" : "Eve"
          }
        ]
      },
      {
        "edges" : [
          {
            "_key" : "11506",
            "_id" : "knows/11506",
            "_from" : "persons/eve",
            "_to" : "persons/alice",
            "_rev" : "_U2FCs2C--B"
          },
          {
            "_key" : "11509",
            "_id" : "knows/11509",
            "_from" : "persons/eve",
            "_to" : "persons/bob",
            "_rev" : "_U2FCs2C--C"
          }
        ],
        "vertices" : [
          {
            "_key" : "alice",
            "_id" : "persons/alice",
            "_rev" : "_U2FCs16--_",
            "name" : "Alice"
          },
          {
            "_key" : "eve",
            "_id" : "persons/eve",
            "_rev" : "_U2FCs2---B",
            "name" : "Eve"
          },
          {
            "_key" : "bob",
            "_id" : "persons/bob",
            "_rev" : "_U2FCs2----",
            "name" : "Bob"
          }
        ]
      },
      {
```

```
      "edges" : [
        {
          "_key" : "11506",
          "_id" : "knows/11506",
          "_from" : "persons/eve",
          "_to" : "persons/alice",
          "_rev" : "_U2FCs2C--B"
        },
        {
          "_key" : "11509",
          "_id" : "knows/11509",
          "_from" : "persons/eve",
          "_to" : "persons/bob",
          "_rev" : "_U2FCs2C--C"
        },
        {
          "_key" : "11500",
          "_id" : "knows/11500",
          "_from" : "persons/bob",
          "_to" : "persons/charlie",
          "_rev" : "_U2FCs2C--_"
        }
      ],
      "vertices" : [
        {
          "_key" : "alice",
          "_id" : "persons/alice",
          "_rev" : "_U2FCs16--_",
          "name" : "Alice"
        },
        {
          "_key" : "eve",
          "_id" : "persons/eve",
          "_rev" : "_U2FCs2---B",
          "name" : "Eve"
        },
        {
          "_key" : "bob",
          "_id" : "persons/bob",
          "_rev" : "_U2FCs2----",
          "name" : "Bob"
        },
        {
          "_key" : "charlie",
          "_id" : "persons/charlie",
          "_rev" : "_U2FCs2---_",
          "name" : "Charlie"
        }
      ]
    },
    {
      "edges" : [
        {
          "_key" : "11506",
          "_id" : "knows/11506",
          "_from" : "persons/eve",
          "_to" : "persons/alice",
          "_rev" : "_U2FCs2C--B"
        },
        {
          "_key" : "11509",
          "_id" : "knows/11509",
          "_from" : "persons/eve",
          "_to" : "persons/bob",
          "_rev" : "_U2FCs2C--C"
        },
        {
          "_key" : "11503",
          "_id" : "knows/11503",
          "_from" : "persons/bob",
          "_to" : "persons/dave",
          "_rev" : "_U2FCs2C--A"
        }
      ],
      "vertices" : [
        {
```

```
            "_key" : "alice",
            "_id" : "persons/alice",
            "_rev" : "_U2FCs16--_",
            "name" : "Alice"
          },
          {
            "_key" : "eve",
            "_id" : "persons/eve",
            "_rev" : "_U2FCs2---B",
            "name" : "Eve"
          },
          {
            "_key" : "bob",
            "_id" : "persons/bob",
            "_rev" : "_U2FCs2----",
            "name" : "Bob"
          },
          {
            "_key" : "dave",
            "_id" : "persons/dave",
            "_rev" : "_U2FCs2---A",
            "name" : "Dave"
          }
        ]
      },
      {
        "edges" : [
          {
            "_key" : "11506",
            "_id" : "knows/11506",
            "_from" : "persons/eve",
            "_to" : "persons/alice",
            "_rev" : "_U2FCs2C--B"
          },
          {
            "_key" : "11509",
            "_id" : "knows/11509",
            "_from" : "persons/eve",
            "_to" : "persons/bob",
            "_rev" : "_U2FCs2C--C"
          },
          {
            "_key" : "11496",
            "_id" : "knows/11496",
            "_from" : "persons/alice",
            "_to" : "persons/bob",
            "_rev" : "_U2FCs2C---"
          }
        ],
        "vertices" : [
          {
            "_key" : "alice",
            "_id" : "persons/alice",
            "_rev" : "_U2FCs16--_",
            "name" : "Alice"
          },
          {
            "_key" : "eve",
            "_id" : "persons/eve",
            "_rev" : "_U2FCs2---B",
            "name" : "Eve"
          },
          {
            "_key" : "bob",
            "_id" : "persons/bob",
            "_rev" : "_U2FCs2----",
            "name" : "Bob"
          },
          {
            "_key" : "alice",
            "_id" : "persons/alice",
            "_rev" : "_U2FCs16--_",
            "name" : "Alice"
          }
        ]
      }
```

```
      ]
    }
  },
  "error" : false,
  "code" : 200
}
```

Example:

Using *postorder* ordering

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/traversal <<EOF
{
  "startVertex" : "persons/alice",
  "graphName" : "knows_graph",
  "direction" : "any",
  "order" : "postorder"
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "result" : {
    "visited" : {
      "vertices" : [
        {
          "_key" : "charlie",
          "_id" : "persons/charlie",
          "_rev" : "_U2FCvfy---",
          "name" : "Charlie"
        },
        {
          "_key" : "dave",
          "_id" : "persons/dave",
          "_rev" : "_U2FCvfy--_",
          "name" : "Dave"
        },
        {
          "_key" : "alice",
          "_id" : "persons/alice",
          "_rev" : "_U2FCvfu---",
          "name" : "Alice"
        },
        {
          "_key" : "eve",
          "_id" : "persons/eve",
          "_rev" : "_U2FCvfy--A",
          "name" : "Eve"
        },
        {
          "_key" : "bob",
          "_id" : "persons/bob",
          "_rev" : "_U2FCvfu--_",
          "name" : "Bob"
        },
        {
          "_key" : "charlie",
          "_id" : "persons/charlie",
          "_rev" : "_U2FCvfy---",
          "name" : "Charlie"
        },
        {
          "_key" : "dave",
          "_id" : "persons/dave",
          "_rev" : "_U2FCvfy--_",
          "name" : "Dave"
        },
        {
          "_key" : "alice",
          "_id" : "persons/alice",
```

```
          "_rev" : "_U2FCvfu---",
          "name" : "Alice"
        },
        {
          "_key" : "bob",
          "_id" : "persons/bob",
          "_rev" : "_U2FCvfu--_",
          "name" : "Bob"
        },
        {
          "_key" : "eve",
          "_id" : "persons/eve",
          "_rev" : "_U2FCvfy--A",
          "name" : "Eve"
        },
        {
          "_key" : "alice",
          "_id" : "persons/alice",
          "_rev" : "_U2FCvfu---",
          "name" : "Alice"
        }
      ],
      "paths" : [
        {
          "edges" : [
            {
              "_key" : "12078",
              "_id" : "knows/12078",
              "_from" : "persons/alice",
              "_to" : "persons/bob",
              "_rev" : "_U2FCvf2---"
            },
            {
              "_key" : "12082",
              "_id" : "knows/12082",
              "_from" : "persons/bob",
              "_to" : "persons/charlie",
              "_rev" : "_U2FCvf2--_"
            }
          ],
          "vertices" : [
            {
              "_key" : "alice",
              "_id" : "persons/alice",
              "_rev" : "_U2FCvfu---",
              "name" : "Alice"
            },
            {
              "_key" : "bob",
              "_id" : "persons/bob",
              "_rev" : "_U2FCvfu--_",
              "name" : "Bob"
            },
            {
              "_key" : "charlie",
              "_id" : "persons/charlie",
              "_rev" : "_U2FCvfy---",
              "name" : "Charlie"
            }
          ]
        },
        {
          "edges" : [
            {
              "_key" : "12078",
              "_id" : "knows/12078",
              "_from" : "persons/alice",
              "_to" : "persons/bob",
              "_rev" : "_U2FCvf2---"
            },
            {
              "_key" : "12085",
              "_id" : "knows/12085",
              "_from" : "persons/bob",
              "_to" : "persons/dave",
              "_rev" : "_U2FCvf6---"
```

```
            }
          ],
          "vertices" : [
            {
              "_key" : "alice",
              "_id" : "persons/alice",
              "_rev" : "_U2FCvfu---",
              "name" : "Alice"
            },
            {
              "_key" : "bob",
              "_id" : "persons/bob",
              "_rev" : "_U2FCvfu--_",
              "name" : "Bob"
            },
            {
              "_key" : "dave",
              "_id" : "persons/dave",
              "_rev" : "_U2FCvfy--_",
              "name" : "Dave"
            }
          ]
        },
        {
          "edges" : [
            {
              "_key" : "12078",
              "_id" : "knows/12078",
              "_from" : "persons/alice",
              "_to" : "persons/bob",
              "_rev" : "_U2FCvf2---"
            },
            {
              "_key" : "12091",
              "_id" : "knows/12091",
              "_from" : "persons/eve",
              "_to" : "persons/bob",
              "_rev" : "_U2FCvg----"
            },
            {
              "_key" : "12088",
              "_id" : "knows/12088",
              "_from" : "persons/eve",
              "_to" : "persons/alice",
              "_rev" : "_U2FCvf6--_"
            }
          ],
          "vertices" : [
            {
              "_key" : "alice",
              "_id" : "persons/alice",
              "_rev" : "_U2FCvfu---",
              "name" : "Alice"
            },
            {
              "_key" : "bob",
              "_id" : "persons/bob",
              "_rev" : "_U2FCvfu--_",
              "name" : "Bob"
            },
            {
              "_key" : "eve",
              "_id" : "persons/eve",
              "_rev" : "_U2FCvfy--A",
              "name" : "Eve"
            },
            {
              "_key" : "alice",
              "_id" : "persons/alice",
              "_rev" : "_U2FCvfu---",
              "name" : "Alice"
            }
          ]
        },
        {
          "edges" : [
```

```
            {
              "_key" : "12078",
              "_id" : "knows/12078",
              "_from" : "persons/alice",
              "_to" : "persons/bob",
              "_rev" : "_U2FCvf2---"
            },
            {
              "_key" : "12091",
              "_id" : "knows/12091",
              "_from" : "persons/eve",
              "_to" : "persons/bob",
              "_rev" : "_U2FCvg----"
            }
          ],
          "vertices" : [
            {
              "_key" : "alice",
              "_id" : "persons/alice",
              "_rev" : "_U2FCvfu---",
              "name" : "Alice"
            },
            {
              "_key" : "bob",
              "_id" : "persons/bob",
              "_rev" : "_U2FCvfu--_",
              "name" : "Bob"
            },
            {
              "_key" : "eve",
              "_id" : "persons/eve",
              "_rev" : "_U2FCvfy--A",
              "name" : "Eve"
            }
          ]
        },
        {
          "edges" : [
            {
              "_key" : "12078",
              "_id" : "knows/12078",
              "_from" : "persons/alice",
              "_to" : "persons/bob",
              "_rev" : "_U2FCvf2---"
            }
          ],
          "vertices" : [
            {
              "_key" : "alice",
              "_id" : "persons/alice",
              "_rev" : "_U2FCvfu---",
              "name" : "Alice"
            },
            {
              "_key" : "bob",
              "_id" : "persons/bob",
              "_rev" : "_U2FCvfu--_",
              "name" : "Bob"
            }
          ]
        },
        {
          "edges" : [
            {
              "_key" : "12088",
              "_id" : "knows/12088",
              "_from" : "persons/eve",
              "_to" : "persons/alice",
              "_rev" : "_U2FCvf6--_"
            },
            {
              "_key" : "12091",
              "_id" : "knows/12091",
              "_from" : "persons/eve",
              "_to" : "persons/bob",
              "_rev" : "_U2FCvg----"
```

```
            },
            {
              "_key" : "12082",
              "_id" : "knows/12082",
              "_from" : "persons/bob",
              "_to" : "persons/charlie",
              "_rev" : "_U2FCvf2--_"
            }
          ],
          "vertices" : [
            {
              "_key" : "alice",
              "_id" : "persons/alice",
              "_rev" : "_U2FCvfu---",
              "name" : "Alice"
            },
            {
              "_key" : "eve",
              "_id" : "persons/eve",
              "_rev" : "_U2FCvfy--A",
              "name" : "Eve"
            },
            {
              "_key" : "bob",
              "_id" : "persons/bob",
              "_rev" : "_U2FCvfu--_",
              "name" : "Bob"
            },
            {
              "_key" : "charlie",
              "_id" : "persons/charlie",
              "_rev" : "_U2FCvfy---",
              "name" : "Charlie"
            }
          ]
        },
        {
          "edges" : [
            {
              "_key" : "12088",
              "_id" : "knows/12088",
              "_from" : "persons/eve",
              "_to" : "persons/alice",
              "_rev" : "_U2FCvf6--_"
            },
            {
              "_key" : "12091",
              "_id" : "knows/12091",
              "_from" : "persons/eve",
              "_to" : "persons/bob",
              "_rev" : "_U2FCvg----"
            },
            {
              "_key" : "12085",
              "_id" : "knows/12085",
              "_from" : "persons/bob",
              "_to" : "persons/dave",
              "_rev" : "_U2FCvf6---"
            }
          ],
          "vertices" : [
            {
              "_key" : "alice",
              "_id" : "persons/alice",
              "_rev" : "_U2FCvfu---",
              "name" : "Alice"
            },
            {
              "_key" : "eve",
              "_id" : "persons/eve",
              "_rev" : "_U2FCvfy--A",
              "name" : "Eve"
            },
            {
              "_key" : "bob",
              "_id" : "persons/bob",
```

```
            "_rev" : "_U2FCvfu--_",
            "name" : "Bob"
          },
          {
            "_key" : "dave",
            "_id" : "persons/dave",
            "_rev" : "_U2FCvfy--_",
            "name" : "Dave"
          }
        ]
      },
      {
        "edges" : [
          {
            "_key" : "12088",
            "_id" : "knows/12088",
            "_from" : "persons/eve",
            "_to" : "persons/alice",
            "_rev" : "_U2FCvf6--_"
          },
          {
            "_key" : "12091",
            "_id" : "knows/12091",
            "_from" : "persons/eve",
            "_to" : "persons/bob",
            "_rev" : "_U2FCvg----"
          },
          {
            "_key" : "12078",
            "_id" : "knows/12078",
            "_from" : "persons/alice",
            "_to" : "persons/bob",
            "_rev" : "_U2FCvf2---"
          }
        ],
        "vertices" : [
          {
            "_key" : "alice",
            "_id" : "persons/alice",
            "_rev" : "_U2FCvfu---",
            "name" : "Alice"
          },
          {
            "_key" : "eve",
            "_id" : "persons/eve",
            "_rev" : "_U2FCvfy--A",
            "name" : "Eve"
          },
          {
            "_key" : "bob",
            "_id" : "persons/bob",
            "_rev" : "_U2FCvfu--_",
            "name" : "Bob"
          },
          {
            "_key" : "alice",
            "_id" : "persons/alice",
            "_rev" : "_U2FCvfu---",
            "name" : "Alice"
          }
        ]
      },
      {
        "edges" : [
          {
            "_key" : "12088",
            "_id" : "knows/12088",
            "_from" : "persons/eve",
            "_to" : "persons/alice",
            "_rev" : "_U2FCvf6--_"
          },
          {
            "_key" : "12091",
            "_id" : "knows/12091",
            "_from" : "persons/eve",
            "_to" : "persons/bob",
```

```
              "_rev" : "_U2FCvg----"
            }
          ],
          "vertices" : [
            {
              "_key" : "alice",
              "_id" : "persons/alice",
              "_rev" : "_U2FCvfu---",
              "name" : "Alice"
            },
            {
              "_key" : "eve",
              "_id" : "persons/eve",
              "_rev" : "_U2FCvfy--A",
              "name" : "Eve"
            },
            {
              "_key" : "bob",
              "_id" : "persons/bob",
              "_rev" : "_U2FCvfu--_",
              "name" : "Bob"
            }
          ]
        },
        {
          "edges" : [
            {
              "_key" : "12088",
              "_id" : "knows/12088",
              "_from" : "persons/eve",
              "_to" : "persons/alice",
              "_rev" : "_U2FCvf6--_"
            }
          ],
          "vertices" : [
            {
              "_key" : "alice",
              "_id" : "persons/alice",
              "_rev" : "_U2FCvfu---",
              "name" : "Alice"
            },
            {
              "_key" : "eve",
              "_id" : "persons/eve",
              "_rev" : "_U2FCvfy--A",
              "name" : "Eve"
            }
          ]
        },
        {
          "edges" : [ ],
          "vertices" : [
            {
              "_key" : "alice",
              "_id" : "persons/alice",
              "_rev" : "_U2FCvfu---",
              "name" : "Alice"
            }
          ]
        }
      ]
    }
  },
  "error" : false,
  "code" : 200
}
```

Example:

Using *backward* item-ordering:

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/traversal <<EOF
```

```
{
  "startVertex" : "persons/alice",
  "graphName" : "knows_graph",
  "direction" : "any",
  "itemOrder" : "backward"
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "result" : {
    "visited" : {
      "vertices" : [
        {
          "_key" : "alice",
          "_id" : "persons/alice",
          "_rev" : "_U2FCsku---",
          "name" : "Alice"
        },
        {
          "_key" : "eve",
          "_id" : "persons/eve",
          "_rev" : "_U2FCsk2---",
          "name" : "Eve"
        },
        {
          "_key" : "bob",
          "_id" : "persons/bob",
          "_rev" : "_U2FCsku--_",
          "name" : "Bob"
        },
        {
          "_key" : "alice",
          "_id" : "persons/alice",
          "_rev" : "_U2FCsku---",
          "name" : "Alice"
        },
        {
          "_key" : "dave",
          "_id" : "persons/dave",
          "_rev" : "_U2FCsky--_",
          "name" : "Dave"
        },
        {
          "_key" : "charlie",
          "_id" : "persons/charlie",
          "_rev" : "_U2FCsky---",
          "name" : "Charlie"
        },
        {
          "_key" : "bob",
          "_id" : "persons/bob",
          "_rev" : "_U2FCsku--_",
          "name" : "Bob"
        },
        {
          "_key" : "eve",
          "_id" : "persons/eve",
          "_rev" : "_U2FCsk2---",
          "name" : "Eve"
        },
        {
          "_key" : "alice",
          "_id" : "persons/alice",
          "_rev" : "_U2FCsku---",
          "name" : "Alice"
        },
        {
          "_key" : "dave",
          "_id" : "persons/dave",
          "_rev" : "_U2FCsky--_",
          "name" : "Dave"
        },
        {
```

```
        "_key" : "charlie",
        "_id" : "persons/charlie",
        "_rev" : "_U2FCsky---",
        "name" : "Charlie"
      }
    ],
    "paths" : [
      {
        "edges" : [ ],
        "vertices" : [
          {
            "_key" : "alice",
            "_id" : "persons/alice",
            "_rev" : "_U2FCsku---",
            "name" : "Alice"
          }
        ]
      },
      {
        "edges" : [
          {
            "_key" : "11421",
            "_id" : "knows/11421",
            "_from" : "persons/eve",
            "_to" : "persons/alice",
            "_rev" : "_U2FCsl----"
          }
        ],
        "vertices" : [
          {
            "_key" : "alice",
            "_id" : "persons/alice",
            "_rev" : "_U2FCsku---",
            "name" : "Alice"
          },
          {
            "_key" : "eve",
            "_id" : "persons/eve",
            "_rev" : "_U2FCsk2---",
            "name" : "Eve"
          }
        ]
      },
      {
        "edges" : [
          {
            "_key" : "11421",
            "_id" : "knows/11421",
            "_from" : "persons/eve",
            "_to" : "persons/alice",
            "_rev" : "_U2FCsl----"
          },
          {
            "_key" : "11424",
            "_id" : "knows/11424",
            "_from" : "persons/eve",
            "_to" : "persons/bob",
            "_rev" : "_U2FCsl---_"
          }
        ],
        "vertices" : [
          {
            "_key" : "alice",
            "_id" : "persons/alice",
            "_rev" : "_U2FCsku---",
            "name" : "Alice"
          },
          {
            "_key" : "eve",
            "_id" : "persons/eve",
            "_rev" : "_U2FCsk2---",
            "name" : "Eve"
          },
          {
            "_key" : "bob",
            "_id" : "persons/bob",
```

```
            "_rev" : "_U2FCsku--_",
            "name" : "Bob"
          }
        ]
      },
      {
        "edges" : [
          {
            "_key" : "11421",
            "_id" : "knows/11421",
            "_from" : "persons/eve",
            "_to" : "persons/alice",
            "_rev" : "_U2FCsl----"
          },
          {
            "_key" : "11424",
            "_id" : "knows/11424",
            "_from" : "persons/eve",
            "_to" : "persons/bob",
            "_rev" : "_U2FCsl---_"
          },
          {
            "_key" : "11411",
            "_id" : "knows/11411",
            "_from" : "persons/alice",
            "_to" : "persons/bob",
            "_rev" : "_U2FCsk2--_"
          }
        ],
        "vertices" : [
          {
            "_key" : "alice",
            "_id" : "persons/alice",
            "_rev" : "_U2FCsku---",
            "name" : "Alice"
          },
          {
            "_key" : "eve",
            "_id" : "persons/eve",
            "_rev" : "_U2FCsk2---",
            "name" : "Eve"
          },
          {
            "_key" : "bob",
            "_id" : "persons/bob",
            "_rev" : "_U2FCsku--_",
            "name" : "Bob"
          },
          {
            "_key" : "alice",
            "_id" : "persons/alice",
            "_rev" : "_U2FCsku---",
            "name" : "Alice"
          }
        ]
      },
      {
        "edges" : [
          {
            "_key" : "11421",
            "_id" : "knows/11421",
            "_from" : "persons/eve",
            "_to" : "persons/alice",
            "_rev" : "_U2FCsl----"
          },
          {
            "_key" : "11424",
            "_id" : "knows/11424",
            "_from" : "persons/eve",
            "_to" : "persons/bob",
            "_rev" : "_U2FCsl---_"
          },
          {
            "_key" : "11418",
            "_id" : "knows/11418",
            "_from" : "persons/bob",
```

```
            "_to" : "persons/dave",
            "_rev" : "_U2FCsk6--_"
          }
        ],
        "vertices" : [
          {
            "_key" : "alice",
            "_id" : "persons/alice",
            "_rev" : "_U2FCsku---",
            "name" : "Alice"
          },
          {
            "_key" : "eve",
            "_id" : "persons/eve",
            "_rev" : "_U2FCsk2---",
            "name" : "Eve"
          },
          {
            "_key" : "bob",
            "_id" : "persons/bob",
            "_rev" : "_U2FCsku--_",
            "name" : "Bob"
          },
          {
            "_key" : "dave",
            "_id" : "persons/dave",
            "_rev" : "_U2FCsky--_",
            "name" : "Dave"
          }
        ]
      },
      {
        "edges" : [
          {
            "_key" : "11421",
            "_id" : "knows/11421",
            "_from" : "persons/eve",
            "_to" : "persons/alice",
            "_rev" : "_U2FCsl----"
          },
          {
            "_key" : "11424",
            "_id" : "knows/11424",
            "_from" : "persons/eve",
            "_to" : "persons/bob",
            "_rev" : "_U2FCsl---_"
          },
          {
            "_key" : "11415",
            "_id" : "knows/11415",
            "_from" : "persons/bob",
            "_to" : "persons/charlie",
            "_rev" : "_U2FCsk6---"
          }
        ],
        "vertices" : [
          {
            "_key" : "alice",
            "_id" : "persons/alice",
            "_rev" : "_U2FCsku---",
            "name" : "Alice"
          },
          {
            "_key" : "eve",
            "_id" : "persons/eve",
            "_rev" : "_U2FCsk2---",
            "name" : "Eve"
          },
          {
            "_key" : "bob",
            "_id" : "persons/bob",
            "_rev" : "_U2FCsku--_",
            "name" : "Bob"
          },
          {
            "_key" : "charlie",
```

```
            "_id" : "persons/charlie",
            "_rev" : "_U2FCsky---",
            "name" : "Charlie"
          }
        ]
      },
      {
        "edges" : [
          {
            "_key" : "11411",
            "_id" : "knows/11411",
            "_from" : "persons/alice",
            "_to" : "persons/bob",
            "_rev" : "_U2FCsk2--_"
          }
        ],
        "vertices" : [
          {
            "_key" : "alice",
            "_id" : "persons/alice",
            "_rev" : "_U2FCsku---",
            "name" : "Alice"
          },
          {
            "_key" : "bob",
            "_id" : "persons/bob",
            "_rev" : "_U2FCsku--_",
            "name" : "Bob"
          }
        ]
      },
      {
        "edges" : [
          {
            "_key" : "11411",
            "_id" : "knows/11411",
            "_from" : "persons/alice",
            "_to" : "persons/bob",
            "_rev" : "_U2FCsk2--_"
          },
          {
            "_key" : "11424",
            "_id" : "knows/11424",
            "_from" : "persons/eve",
            "_to" : "persons/bob",
            "_rev" : "_U2FCsl---_"
          }
        ],
        "vertices" : [
          {
            "_key" : "alice",
            "_id" : "persons/alice",
            "_rev" : "_U2FCsku---",
            "name" : "Alice"
          },
          {
            "_key" : "bob",
            "_id" : "persons/bob",
            "_rev" : "_U2FCsku--_",
            "name" : "Bob"
          },
          {
            "_key" : "eve",
            "_id" : "persons/eve",
            "_rev" : "_U2FCsk2---",
            "name" : "Eve"
          }
        ]
      },
      {
        "edges" : [
          {
            "_key" : "11411",
            "_id" : "knows/11411",
            "_from" : "persons/alice",
            "_to" : "persons/bob",
```

```
            "_rev" : "_U2FCsk2--_"
          },
          {
            "_key" : "11424",
            "_id" : "knows/11424",
            "_from" : "persons/eve",
            "_to" : "persons/bob",
            "_rev" : "_U2FCsl---_"
          },
          {
            "_key" : "11421",
            "_id" : "knows/11421",
            "_from" : "persons/eve",
            "_to" : "persons/alice",
            "_rev" : "_U2FCsl----"
          }
        ],
        "vertices" : [
          {
            "_key" : "alice",
            "_id" : "persons/alice",
            "_rev" : "_U2FCsku---",
            "name" : "Alice"
          },
          {
            "_key" : "bob",
            "_id" : "persons/bob",
            "_rev" : "_U2FCsku--_",
            "name" : "Bob"
          },
          {
            "_key" : "eve",
            "_id" : "persons/eve",
            "_rev" : "_U2FCsk2---",
            "name" : "Eve"
          },
          {
            "_key" : "alice",
            "_id" : "persons/alice",
            "_rev" : "_U2FCsku---",
            "name" : "Alice"
          }
        ]
      },
      {
        "edges" : [
          {
            "_key" : "11411",
            "_id" : "knows/11411",
            "_from" : "persons/alice",
            "_to" : "persons/bob",
            "_rev" : "_U2FCsk2--_"
          },
          {
            "_key" : "11418",
            "_id" : "knows/11418",
            "_from" : "persons/bob",
            "_to" : "persons/dave",
            "_rev" : "_U2FCsk6--_"
          }
        ],
        "vertices" : [
          {
            "_key" : "alice",
            "_id" : "persons/alice",
            "_rev" : "_U2FCsku---",
            "name" : "Alice"
          },
          {
            "_key" : "bob",
            "_id" : "persons/bob",
            "_rev" : "_U2FCsku--_",
            "name" : "Bob"
          },
          {
            "_key" : "dave",
```

```
              "_id" : "persons/dave",
              "_rev" : "_U2FCsky--_",
              "name" : "Dave"
            }
          ]
        },
        {
          "edges" : [
            {
              "_key" : "11411",
              "_id" : "knows/11411",
              "_from" : "persons/alice",
              "_to" : "persons/bob",
              "_rev" : "_U2FCsk2--_"
            },
            {
              "_key" : "11415",
              "_id" : "knows/11415",
              "_from" : "persons/bob",
              "_to" : "persons/charlie",
              "_rev" : "_U2FCsk6---"
            }
          ],
          "vertices" : [
            {
              "_key" : "alice",
              "_id" : "persons/alice",
              "_rev" : "_U2FCsku---",
              "name" : "Alice"
            },
            {
              "_key" : "bob",
              "_id" : "persons/bob",
              "_rev" : "_U2FCsku--_",
              "name" : "Bob"
            },
            {
              "_key" : "charlie",
              "_id" : "persons/charlie",
              "_rev" : "_U2FCsky---",
              "name" : "Charlie"
            }
          ]
        }
      ]
    }
  },
  "error" : false,
  "code" : 200
}
```

Example:

Edges should only be included once globally, but nodes are included every time they are visited

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/traversal <<EOF
{
  "startVertex" : "persons/alice",
  "graphName" : "knows_graph",
  "direction" : "any",
  "uniqueness" : {
    "vertices" : "none",
    "edges" : "global"
  }
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "result" : {
```

```
    "visited" : {
      "vertices" : [
        {
          "_key" : "alice",
          "_id" : "persons/alice",
          "_rev" : "_U2FCtGS---",
          "name" : "Alice"
        },
        {
          "_key" : "bob",
          "_id" : "persons/bob",
          "_rev" : "_U2FCtGS--_",
          "name" : "Bob"
        },
        {
          "_key" : "charlie",
          "_id" : "persons/charlie",
          "_rev" : "_U2FCtGW---",
          "name" : "Charlie"
        },
        {
          "_key" : "dave",
          "_id" : "persons/dave",
          "_rev" : "_U2FCtGW--_",
          "name" : "Dave"
        },
        {
          "_key" : "eve",
          "_id" : "persons/eve",
          "_rev" : "_U2FCtGa---",
          "name" : "Eve"
        },
        {
          "_key" : "alice",
          "_id" : "persons/alice",
          "_rev" : "_U2FCtGS---",
          "name" : "Alice"
        }
      ],
      "paths" : [
        {
          "edges" : [ ],
          "vertices" : [
            {
              "_key" : "alice",
              "_id" : "persons/alice",
              "_rev" : "_U2FCtGS---",
              "name" : "Alice"
            }
          ]
        },
        {
          "edges" : [
            {
              "_key" : "11581",
              "_id" : "knows/11581",
              "_from" : "persons/alice",
              "_to" : "persons/bob",
              "_rev" : "_U2FCtGa--_"
            }
          ],
          "vertices" : [
            {
              "_key" : "alice",
              "_id" : "persons/alice",
              "_rev" : "_U2FCtGS---",
              "name" : "Alice"
            },
            {
              "_key" : "bob",
              "_id" : "persons/bob",
              "_rev" : "_U2FCtGS--_",
              "name" : "Bob"
            }
          ]
        },
```

```
    {
      "edges" : [
        {
          "_key" : "11581",
          "_id" : "knows/11581",
          "_from" : "persons/alice",
          "_to" : "persons/bob",
          "_rev" : "_U2FCtGa--_"
        },
        {
          "_key" : "11585",
          "_id" : "knows/11585",
          "_from" : "persons/bob",
          "_to" : "persons/charlie",
          "_rev" : "_U2FCtGe---"
        }
      ],
      "vertices" : [
        {
          "_key" : "alice",
          "_id" : "persons/alice",
          "_rev" : "_U2FCtGS---",
          "name" : "Alice"
        },
        {
          "_key" : "bob",
          "_id" : "persons/bob",
          "_rev" : "_U2FCtGS--_",
          "name" : "Bob"
        },
        {
          "_key" : "charlie",
          "_id" : "persons/charlie",
          "_rev" : "_U2FCtGW---",
          "name" : "Charlie"
        }
      ]
    },
    {
      "edges" : [
        {
          "_key" : "11581",
          "_id" : "knows/11581",
          "_from" : "persons/alice",
          "_to" : "persons/bob",
          "_rev" : "_U2FCtGa--_"
        },
        {
          "_key" : "11588",
          "_id" : "knows/11588",
          "_from" : "persons/bob",
          "_to" : "persons/dave",
          "_rev" : "_U2FCtGe--_"
        }
      ],
      "vertices" : [
        {
          "_key" : "alice",
          "_id" : "persons/alice",
          "_rev" : "_U2FCtGS---",
          "name" : "Alice"
        },
        {
          "_key" : "bob",
          "_id" : "persons/bob",
          "_rev" : "_U2FCtGS--_",
          "name" : "Bob"
        },
        {
          "_key" : "dave",
          "_id" : "persons/dave",
          "_rev" : "_U2FCtGW--_",
          "name" : "Dave"
        }
      ]
    },
```

```
        {
          "edges" : [
            {
              "_key" : "11581",
              "_id" : "knows/11581",
              "_from" : "persons/alice",
              "_to" : "persons/bob",
              "_rev" : "_U2FCtGa--_"
            },
            {
              "_key" : "11594",
              "_id" : "knows/11594",
              "_from" : "persons/eve",
              "_to" : "persons/bob",
              "_rev" : "_U2FCtGi---"
            }
          ],
          "vertices" : [
            {
              "_key" : "alice",
              "_id" : "persons/alice",
              "_rev" : "_U2FCtGS---",
              "name" : "Alice"
            },
            {
              "_key" : "bob",
              "_id" : "persons/bob",
              "_rev" : "_U2FCtGS--_",
              "name" : "Bob"
            },
            {
              "_key" : "eve",
              "_id" : "persons/eve",
              "_rev" : "_U2FCtGa---",
              "name" : "Eve"
            }
          ]
        },
        {
          "edges" : [
            {
              "_key" : "11581",
              "_id" : "knows/11581",
              "_from" : "persons/alice",
              "_to" : "persons/bob",
              "_rev" : "_U2FCtGa--_"
            },
            {
              "_key" : "11594",
              "_id" : "knows/11594",
              "_from" : "persons/eve",
              "_to" : "persons/bob",
              "_rev" : "_U2FCtGi---"
            },
            {
              "_key" : "11591",
              "_id" : "knows/11591",
              "_from" : "persons/eve",
              "_to" : "persons/alice",
              "_rev" : "_U2FCtGe--A"
            }
          ],
          "vertices" : [
            {
              "_key" : "alice",
              "_id" : "persons/alice",
              "_rev" : "_U2FCtGS---",
              "name" : "Alice"
            },
            {
              "_key" : "bob",
              "_id" : "persons/bob",
              "_rev" : "_U2FCtGS--_",
              "name" : "Bob"
            },
            {
```

```
              "_key" : "eve",
              "_id" : "persons/eve",
              "_rev" : "_U2FCtGa---",
              "name" : "Eve"
            },
            {
              "_key" : "alice",
              "_id" : "persons/alice",
              "_rev" : "_U2FCtGS---",
              "name" : "Alice"
            }
          ]
        }
      ]
    }
  },
  "error" : false,
  "code" : 200
}
```

Example:

If the underlying graph is cyclic, *maxIterations* should be set

The underlying graph has two vertices *Alice* and *Bob*. With the directed edges:

- *Alice* knows *Bob*
- *Bob* knows *Alice*

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/traversal <<EOF
{
  "startVertex" : "persons/alice",
  "graphName" : "knows_graph",
  "direction" : "any",
  "uniqueness" : {
    "vertices" : "none",
    "edges" : "none"
  },
  "maxIterations" : 5
}
EOF

HTTP/1.1 500 Internal Server Error
content-type: application/json; charset=utf-8

{
  "error" : true,
  "code" : 500,
  "errorNum" : 1909,
  "errorMessage" : "too many iterations - try increasing the value of 'maxIterations'"
}
```

Return Codes

- *200*: If the traversal is fully executed *HTTP 200* will be returned.
- *400*: If the traversal specification is either missing or malformed, the server will respond with *HTTP 400*.
- *404*: The server will responded with *HTTP 404* if the specified edge collection does not exist, or the specified start vertex cannot be found.
- *500*: The server will responded with *HTTP 500* when an error occurs inside the traversal or if a traversal performs more than *maxIterations* iterations.

Examples

In the following examples the underlying graph will contain five persons *Alice*, *Bob*, *Charlie*, *Dave* and *Eve*. We will have the following directed relations:

- *Alice* knows *Bob*
- *Bob* knows *Charlie*

- *Bob* knows *Dave*
- *Eve* knows *Alice*
- *Eve* knows *Bob* The starting vertex will always be Alice. Follow only outbound edges

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/traversal <<EOF
{
  "startVertex" : "persons/alice",
  "graphName" : "knows_graph",
  "direction" : "outbound"
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Follow only inbound edges

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/traversal <<EOF
{
  "startVertex" : "persons/alice",
  "graphName" : "knows_graph",
  "direction" : "inbound"
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Follow any direction of edges

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/traversal <<EOF
{
  "startVertex" : "persons/alice",
  "graphName" : "knows_graph",
  "direction" : "any",
  "uniqueness" : {
    "vertices" : "none",
    "edges" : "global"
  }
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Excluding *Charlie* and *Bob*

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/traversal <<EOF
{
  "startVertex" : "persons/alice",
  "graphName" : "knows_graph",
  "direction" : "outbound",
  "filter" : "if (vertex.name === \"Bob\" ||     vertex.name === \"Charlie\") {  return
\"exclude\";}return;"
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Do not follow edges from *Bob*

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/traversal <<EOF
{
  "startVertex" : "persons/alice",
  "graphName" : "knows_graph",
  "direction" : "outbound",
  "filter" : "if (vertex.name === \"Bob\") {return \"prune\";}return;"
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Visit only nodes in a depth of at least 2

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/traversal <<EOF
{
  "startVertex" : "persons/alice",
  "graphName" : "knows_graph",
  "direction" : "outbound",
  "minDepth" : 2
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Visit only nodes in a depth of at most 1

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/traversal <<EOF
{
  "startVertex" : "persons/alice",
  "graphName" : "knows_graph",
  "direction" : "outbound",
  "maxDepth" : 1
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Using a visitor function to return vertex ids only

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/traversal <<EOF
{
  "startVertex" : "persons/alice",
  "graphName" : "knows_graph",
  "direction" : "outbound",
  "visitor" : "result.visited.vertices.push(vertex._id);"
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Count all visited nodes and return a list of nodes only

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/traversal <<EOF
{
  "startVertex" : "persons/alice",
  "graphName" : "knows_graph",
  "direction" : "outbound",
  "init" : "result.visited = 0; result.myVertices = [ ];",
  "visitor" : "result.visited++; result.myVertices.push(vertex);"
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Expand only inbound edges of *Alice* and outbound edges of *Eve*

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/traversal <<EOF
{
  "startVertex" : "persons/alice",
  "graphName" : "knows_graph",
  "expander" : "var connections = [ ];if (vertex.name === \"Alice\")
{config.datasource.getInEdges(vertex).forEach(function (e) {connections.push({ vertex:
require(\"internal\").db._document(e._from), edge: e});});}if (vertex.name === \"Eve\")
{config.datasource.getOutEdges(vertex).forEach(function (e) {connections.push({vertex:
require(\"internal\").db._document(e._to), edge: e});});}return connections;"
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Follow the *depthfirst* strategy

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/traversal <<EOF
{
  "startVertex" : "persons/alice",
  "graphName" : "knows_graph",
  "direction" : "any",
  "strategy" : "depthfirst"
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Using *postorder* ordering

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/traversal <<EOF
{
  "startVertex" : "persons/alice",
  "graphName" : "knows_graph",
  "direction" : "any",
  "order" : "postorder"
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Using *backward* item-ordering:

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/traversal <<EOF
{
  "startVertex" : "persons/alice",
  "graphName" : "knows_graph",
  "direction" : "any",
  "itemOrder" : "backward"
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Edges should only be included once globally, but nodes are included every time they are visited

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/traversal <<EOF
{
  "startVertex" : "persons/alice",
  "graphName" : "knows_graph",
  "direction" : "any",
  "uniqueness" : {
    "vertices" : "none",
    "edges" : "global"
  }
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

If the underlying graph is cyclic, *maxIterations* should be set The underlying graph has two vertices *Alice* and *Bob*. With the directed edges:

- *Alice* knows *Bob*
- *Bob* knows *Alice*

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/traversal <<EOF
{
  "startVertex" : "persons/alice",
  "graphName" : "knows_graph",
  "direction" : "any",
  "uniqueness" : {
    "vertices" : "none",
    "edges" : "none"
  },
  "maxIterations" : 5
}
EOF

HTTP/1.1 500 Internal Server Error
content-type: application/json; charset=utf-8
```

show response body

All examples were using this graph:

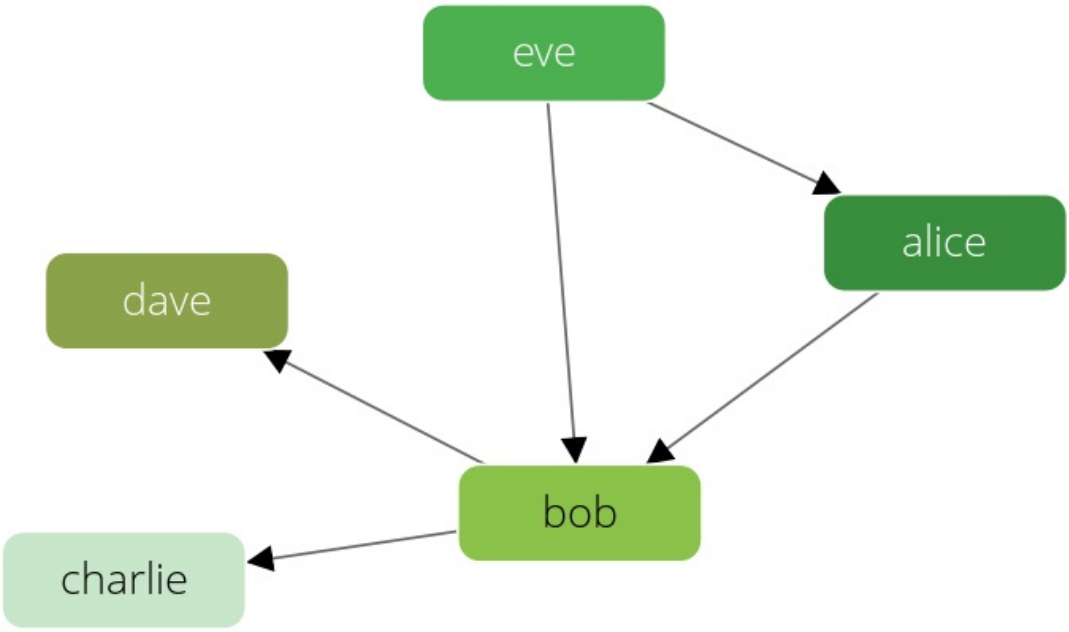

HTTP Interface for AQL Query Cursors

Database Cursors

This is an introduction to ArangoDB's HTTP Interface for Queries. Results of AQL and simple queries are returned as cursors in order to batch the communication between server and client. Each call returns a number of documents in a batch and an indication if the current batch has been the final batch. Depending on the query, the total number of documents in the result set might or might not be known in advance. In order to free server resources the client should delete the cursor as soon as it is no longer needed.

To execute a query, the query details need to be shipped from the client to the server via an HTTP POST request.

Retrieving query results

Select queries are executed on-the-fly on the server and the result set will be returned back to the client.

There are two ways the client can get the result set from the server:

- In a single roundtrip
- Using a cursor

Single roundtrip

The server will only transfer a certain number of result documents back to the client in one roundtrip. This number is controllable by the client by setting the *batchSize* attribute when issuing the query.

If the complete result can be transferred to the client in one go, the client does not need to issue any further request. The client can check whether it has retrieved the complete result set by checking the *hasMore* attribute of the result set. If it is set to *false*, then the client has fetched the complete result set from the server. In this case no server side cursor will be created.

```
> curl --data @- -X POST --dump - http://localhost:8529/_api/cursor
{ "query" : "FOR u IN users LIMIT 2 RETURN u", "count" : true, "batchSize" : 2 }

HTTP/1.1 201 Created
Content-type: application/json

{
  "hasMore" : false,
  "error" : false,
  "result" : [
    {
      "name" : "user1",
      "_rev" : "210304551",
      "_key" : "210304551",
      "_id" : "users/210304551"
    },
    {
      "name" : "user2",
      "_rev" : "210304552",
      "_key" : "210304552",
      "_id" : "users/210304552"
    }
  ],
  "code" : 201,
  "count" : 2
}
```

Using a cursor

If the result set contains more documents than should be transferred in a single roundtrip (i.e. as set via the *batchSize* attribute), the server will return the first few documents and create a temporary cursor. The cursor identifier will also be returned to the client. The server will put the cursor identifier in the *id* attribute of the response object. Furthermore, the *hasMore* attribute of the response object will be set to *true*. This is an indication for the client that there are additional results to fetch from the server.

Examples:

Create and extract first batch:

```
> curl --data @- -X POST --dump - http://localhost:8529/_api/cursor
{ "query" : "FOR u IN users LIMIT 5 RETURN u", "count" : true, "batchSize" : 2 }

HTTP/1.1 201 Created
Content-type: application/json

{
  "hasMore" : true,
  "error" : false,
  "id" : "26011191",
  "result" : [
    {
      "name" : "user1",
      "_rev" : "258801191",
      "_key" : "258801191",
      "_id" : "users/258801191"
    },
    {
      "name" : "user2",
      "_rev" : "258801192",
      "_key" : "258801192",
      "_id" : "users/258801192"
    }
  ],
  "code" : 201,
  "count" : 5
}
```

Extract next batch, still have more:

```
> curl -X PUT --dump - http://localhost:8529/_api/cursor/26011191

HTTP/1.1 200 OK
Content-type: application/json

{
  "hasMore" : true,
  "error" : false,
  "id" : "26011191",
  "result": [
    {
      "name" : "user3",
      "_rev" : "258801193",
      "_key" : "258801193",
      "_id" : "users/258801193"
    },
    {
      "name" : "user4",
      "_rev" : "258801194",
      "_key" : "258801194",
      "_id" : "users/258801194"
    }
  ],
  "code" : 200,
  "count" : 5
}
```

Extract next batch, done:

```
> curl -X PUT --dump - http://localhost:8529/_api/cursor/26011191

HTTP/1.1 200 OK
Content-type: application/json

{
  "hasMore" : false,
  "error" : false,
  "result" : [
    {
      "name" : "user5",
      "_rev" : "258801195",
      "_key" : "258801195",
      "_id" : "users/258801195"
    }
  ],
  "code" : 200,
  "count" : 5
}
```

Do not do this because *hasMore* now has a value of false:

```
> curl -X PUT --dump - http://localhost:8529/_api/cursor/26011191

HTTP/1.1 404 Not Found
Content-type: application/json

{
  "errorNum": 1600,
  "errorMessage": "cursor not found: disposed or unknown cursor",
  "error": true,
  "code": 404
}
```

Modifying documents

The `_api/cursor` endpoint can also be used to execute modifying queries.

The following example appends a value into the array `arrayValue` of the document with key `test` in the collection `documents`. Normal update behavior is to replace the attribute completely, and using an update AQL query with the `PUSH()` function allows to append to the array.

```
curl --data @- -X POST --dump http://127.0.0.1:8529/_api/cursor
{ "query": "FOR doc IN documents FILTER doc._key == @myKey UPDATE doc._key WITH { arrayValue: PUSH(doc.arrayValue, @v
alue) } IN documents","bindVars": { "myKey": "test", "value": 42 } }

HTTP/1.1 201 Created
Content-type: application/json; charset=utf-8

{
  "result" : [],
  "hasMore" : false,
  "extra" : {
    "stats" : {
      "writesExecuted" : 1,
      "writesIgnored" : 0,
      "scannedFull" : 0,
      "scannedIndex" : 1,
      "filtered" : 0
    },
    "warnings" : []
  },
  "error" : false,
  "code" : 201
}
```

Setting a memory limit

To set a memory limit for the query, the *memoryLimit* option can be passed to the server. The memory limit specifies the maximum number of bytes that the query is allowed to use. When a single AQL query reaches the specified limit value, the query will be aborted with a *resource limit exceeded* exception. In a cluster, the memory accounting is done per shard, so the limit value is effectively a memory limit per query per shard.

```
> curl --data @- -X POST --dump - http://localhost:8529/_api/cursor
{ "query" : "FOR i IN 1..100000 SORT i RETURN i", "memoryLimit" : 100000 }

HTTP/1.1 500 Internal Server Error
Server: ArangoDB
Connection: Keep-Alive
Content-Type: application/json; charset=utf-8
Content-Length: 115

{"error":true,"errorMessage":"query would use more memory than allowed (while executing)","code":500,"errorNum":32}
```

If no memory limit is specified, then the server default value (controlled by startup option *--query.memory-limit* will be used for restricting the maximum amount of memory the query can use. A memory limit value of *0* means that the maximum amount of memory for the query is not restricted.

Accessing Cursors via HTTP

Create cursor

create a cursor and return the first results

`POST /_api/cursor` A JSON object describing the query and query parameters.

A json post document with these Properties is required:

- **count**: indicates whether the number of documents in the result set should be returned in the "count" attribute of the result. Calculating the "count" attribute might have a performance impact for some queries in the future so this option is turned off by default, and "count" is only returned when requested.
- **batchSize**: maximum number of result documents to be transferred from the server to the client in one roundtrip. If this attribute is not set, a server-controlled default value will be used. A *batchSize* value of *0* is disallowed.
- **cache**: flag to determine whether the AQL query cache shall be used. If set to *false*, then any query cache lookup will be skipped for the query. If set to *true*, it will lead to the query cache being checked for the query if the query cache mode is either *on* or *demand*.
- **memoryLimit**: the maximum number of memory (measured in bytes) that the query is allowed to use. If set, then the query will fail with error "resource limit exceeded" in case it allocates too much memory. A value of *0* indicates that there is no memory limit.
- **ttl**: The time-to-live for the cursor (in seconds). The cursor will be removed on the server automatically after the specified amount of time. This is useful to ensure garbage collection of cursors that are not fully fetched by clients. If not set, a server-defined value will be used.
- **query**: contains the query string to be executed
- **bindVars** (object): key/value pairs representing the bind parameters.
- **options**:
  - **profile**: If set to *true*, then the additional query profiling information will be returned in the sub-attribute *profile* of the *extra* return attribute if the query result is not served from the query cache.
  - **optimizer.rules** (string): A list of to-be-included or to-be-excluded optimizer rules can be put into this attribute, telling the optimizer to include or exclude specific rules. To disable a rule, prefix its name with a `-`, to enable a rule, prefix it with a `+`. There is also a pseudo-rule `all`, which will match all optimizer rules.
  - **satelliteSyncWait**: This *enterprise* parameter allows to configure how long a DBServer will have time to bring the satellite collections involved in the query into sync. The default value is *60.0* (seconds). When the max time has been reached the query will be stopped.
  - **fullCount**: if set to *true* and the query contains a *LIMIT* clause, then the result will have an *extra* attribute with the sub-attributes *stats* and *fullCount*, `{ ... , "extra": { "stats": { "fullCount": 123 } } }`. The *fullCount* attribute will contain the number of documents in the result before the last LIMIT in the query was applied. It can be used to count the number of documents that match certain filter criteria, but only return a subset of them, in one go. It is thus similar to MySQL's *SQL_CALC_FOUND_ROWS* hint. Note that setting the option will disable a few LIMIT optimizations and may lead to more documents being processed, and thus make queries run longer. Note that the *fullCount* attribute will only be present in the result if the query has a LIMIT clause and the LIMIT clause is actually used in the query.
  - **maxPlans**: Limits the maximum number of plans that are created by the AQL query optimizer.

The query details include the query string plus optional query options and bind parameters. These values need to be passed in a JSON representation in the body of the POST request.

A json document with these Properties is returned:

HTTP 201

is returned if the result set can be created by the server.

- **count**: the total number of result documents available (only available if the query was executed with the *count* attribute set)
- **code**: the HTTP status code
- **extra**: an optional JSON object with extra information about the query result contained in its *stats* sub-attribute. For data-modification queries, the *extra.stats* sub-attribute will contain the number of modified documents and the number of documents that could not be modified due to an error (if *ignoreErrors* query option is specified)

- **cached**: a boolean flag indicating whether the query result was served from the query cache or not. If the query result is served from the query cache, the *extra* return attribute will not contain any *stats* sub-attribute and no *profile* sub-attribute.
- **hasMore**: A boolean indicator whether there are more results available for the cursor on the server
- **result** (anonymous json object): an array of result documents (might be empty if query has no results)
- **error**: A flag to indicate that an error occurred (*false* in this case)
- **id**: id of temporary cursor created on the server (optional, see above)

A json document with these Properties is returned:

HTTP 400

is returned if the JSON representation is malformed or the query specification is missing from the request. If the JSON representation is malformed or the query specification is missing from the request, the server will respond with *HTTP 400*. The body of the response will contain a JSON object with additional error details. The object has the following attributes:

- **errorMessage**: a descriptive error message If the query specification is complete, the server will process the query. If an error occurs during query processing, the server will respond with *HTTP 400*. Again, the body of the response will contain details about the error. A list of query errors can be found here.
- **errorNum**: the server error number
- **code**: the HTTP status code
- **error**: boolean flag to indicate that an error occurred (*true* in this case)

Example:

Execute a query and extract the result in a single go

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/cursor <<EOF
{
  "query" : "FOR p IN products LIMIT 2 RETURN p",
  "count" : true,
  "batchSize" : 2
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8

{
  "result" : [
    {
      "_key" : "9566",
      "_id" : "products/9566",
      "_rev" : "_U2FCBkW---",
      "hello1" : "world1"
    },
    {
      "_key" : "9570",
      "_id" : "products/9570",
      "_rev" : "_U2FCBka---",
      "hello2" : "world1"
    }
  ],
  "hasMore" : false,
  "count" : 2,
  "cached" : false,
  "extra" : {
    "stats" : {
      "writesExecuted" : 0,
      "writesIgnored" : 0,
      "scannedFull" : 2,
      "scannedIndex" : 0,
      "filtered" : 0,
      "executionTime" : 0.0001418590545654297
    },
    "warnings" : [ ]
  },
  "error" : false,
  "code" : 201
}
```

Example:

Execute a query and extract a part of the result

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/cursor <<EOF
{
  "query" : "FOR p IN products LIMIT 5 RETURN p",
  "count" : true,
  "batchSize" : 2
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8

{
  "error" : false,
  "code" : 201,
  "result" : [
    {
      "_key" : "9550",
      "_id" : "products/9550",
      "_rev" : "_U2FCBcm---",
      "hello5" : "world1"
    },
    {
      "_key" : "9543",
      "_id" : "products/9543",
      "_rev" : "_U2FCBci--A",
      "hello3" : "world1"
    }
  ],
  "hasMore" : true,
  "id" : "9553",
  "count" : 5,
  "extra" : {
    "stats" : {
      "writesExecuted" : 0,
      "writesIgnored" : 0,
      "scannedFull" : 5,
      "scannedIndex" : 0,
      "filtered" : 0,
      "executionTime" : 0.0001068115234375
    },
    "warnings" : [ ]
  },
  "cached" : false
}
```

Example:

Using the query option "fullCount"

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/cursor <<EOF
{
  "query" : "FOR i IN 1..1000 FILTER i > 500 LIMIT 10 RETURN i",
  "count" : true,
  "options" : {
    "fullCount" : true
  }
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8

{
  "result" : [
    501,
    502,
    503,
    504,
    505,
    506,
    507,
    508,
    509,
    510
  ],
  "hasMore" : false,
  "count" : 10,
  "cached" : false,
  "extra" : {
    "stats" : {
      "writesExecuted" : 0,
      "writesIgnored" : 0,
      "scannedFull" : 0,
      "scannedIndex" : 0,
      "filtered" : 500,
      "fullCount" : 500,
      "executionTime" : 0.0004448890686035156
    },
    "warnings" : [ ]
  },
  "error" : false,
  "code" : 201
}
```

Example:

Enabling and disabling optimizer rules

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/cursor <<EOF
{
  "query" : "FOR i IN 1..10 LET a = 1 LET b = 2 FILTER a + b == 3 RETURN i",
  "count" : true,
  "options" : {
    "maxPlans" : 1,
    "optimizer" : {
      "rules" : [
        "-all",
        "+remove-unnecessary-filters"
      ]
    }
  }
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8

{
  "result" : [
    1,
    2,
    3,
    4,
    5,
    6,
    7,
    8,
    9,
    10
  ],
  "hasMore" : false,
  "count" : 10,
  "cached" : false,
  "extra" : {
    "stats" : {
      "writesExecuted" : 0,
      "writesIgnored" : 0,
      "scannedFull" : 0,
      "scannedIndex" : 0,
      "filtered" : 0,
      "executionTime" : 0.0013229846954345703
    },
    "warnings" : [ ]
  },
  "error" : false,
  "code" : 201
}
```

Example:

Execute a data-modification query and retrieve the number of modified documents

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/cursor <<EOF
{
  "query" : "FOR p IN products REMOVE p IN products"
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8

{
  "result" : [ ],
  "hasMore" : false,
  "cached" : false,
  "extra" : {
    "stats" : {
      "writesExecuted" : 2,
      "writesIgnored" : 0,
      "scannedFull" : 2,
      "scannedIndex" : 0,
      "filtered" : 0,
      "executionTime" : 0.0002110004425048828
    },
    "warnings" : [ ]
  },
  "error" : false,
  "code" : 201
}
```

Example:

Execute a data-modification query with option *ignoreErrors*

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/cursor <<EOF
{
  "query" : "REMOVE 'bar' IN products OPTIONS { ignoreErrors: true }"
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8

{
  "result" : [ ],
  "hasMore" : false,
  "cached" : false,
  "extra" : {
    "stats" : {
      "writesExecuted" : 0,
      "writesIgnored" : 1,
      "scannedFull" : 0,
      "scannedIndex" : 0,
      "filtered" : 0,
      "executionTime" : 0.00020003318786621094
    },
    "warnings" : [ ]
  },
  "error" : false,
  "code" : 201
}
```

Example:

Bad query - Missing body

```
shell> curl -X POST --dump - http://localhost:8529/_api/cursor

HTTP/1.1 400 Bad Request
content-type: application/json; charset=utf-8

{
  "error" : true,
  "errorMessage" : "query is empty",
  "code" : 400,
  "errorNum" : 1502
}
```

Example:

Bad query - Unknown collection

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/cursor <<EOF
{
  "query" : "FOR u IN unknowncoll LIMIT 2 RETURN u",
  "count" : true,
  "batchSize" : 2
}
EOF

HTTP/1.1 404 Not Found
content-type: application/json; charset=utf-8

{
  "error" : true,
  "errorMessage" : "collection not found (unknowncoll)",
  "code" : 404,
  "errorNum" : 1203
}
```

Example:

Bad query - Execute a data-modification query that attempts to remove a non-existing document

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/cursor <<EOF
{
  "query" : "REMOVE 'foo' IN products"
}
EOF

HTTP/1.1 404 Not Found
content-type: application/json; charset=utf-8

{
  "error" : true,
  "errorMessage" : "document not found (while executing)",
  "code" : 404,
  "errorNum" : 1202
}
```

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/cursor <<EOF
{
  "query" : "FOR p IN products LIMIT 2 RETURN p",
  "count" : true,
  "batchSize" : 2
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8
```

show response body

Execute a query and extract a part of the result

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/cursor <<EOF
{
  "query" : "FOR p IN products LIMIT 5 RETURN p",
  "count" : true,
  "batchSize" : 2
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8
```

show response body

Using the query option "fullCount"

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/cursor <<EOF
{
  "query" : "FOR i IN 1..1000 FILTER i > 500 LIMIT 10 RETURN i",
  "count" : true,
  "options" : {
    "fullCount" : true
  }
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8
```

show response body

Enabling and disabling optimizer rules

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/cursor <<EOF
{
  "query" : "FOR i IN 1..10 LET a = 1 LET b = 2 FILTER a + b == 3 RETURN i",
  "count" : true,
  "options" : {
    "maxPlans" : 1,
    "optimizer" : {
      "rules" : [
        "-all",
        "+remove-unnecessary-filters"
      ]
    }
  }
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8
```

show response body

Execute a data-modification query and retrieve the number of modified documents

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/cursor <<EOF
{
  "query" : "FOR p IN products REMOVE p IN products"
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8
```

show response body

Execute a data-modification query with option *ignoreErrors*

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/cursor <<EOF
{
  "query" : "REMOVE 'bar' IN products OPTIONS { ignoreErrors: true }"
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8
```

show response body

Bad query - Missing body

```
shell> curl -X POST --dump - http://localhost:8529/_api/cursor

HTTP/1.1 400 Bad Request
content-type: application/json; charset=utf-8
```

show response body

Bad query - Unknown collection

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/cursor <<EOF
{
  "query" : "FOR u IN unknowncoll LIMIT 2 RETURN u",
  "count" : true,
  "batchSize" : 2
}
EOF

HTTP/1.1 404 Not Found
content-type: application/json; charset=utf-8
```

show response body

Bad query - Execute a data-modification query that attempts to remove a non-existing document

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/cursor <<EOF
{
  "query" : "REMOVE 'foo' IN products"
}
EOF

HTTP/1.1 404 Not Found
content-type: application/json; charset=utf-8
```

show response body

Read next batch from cursor

return the next results from an existing cursor

`PUT /_api/cursor/{cursor-identifier}`

Path Parameters

- *cursor-identifier* (required): The name of the cursor

If the cursor is still alive, returns an object with the following attributes:

- *id*: the *cursor-identifier*
- *result*: a list of documents for the current batch
- *hasMore*: *false* if this was the last batch
- *count*: if present the total number of elements

Note that even if *hasMore* returns *true*, the next call might still return no documents. If, however, *hasMore* is *false*, then the cursor is exhausted. Once the *hasMore* attribute has a value of *false*, the client can stop.

Example:

Valid request for next batch

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/cursor <<EOF
{
  "query" : "FOR p IN products LIMIT 5 RETURN p",
  "count" : true,
  "batchSize" : 2
}
EOF

shell> curl -X PUT --dump - http://localhost:8529/_api/cursor/9645

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "error" : false,
  "code" : 200,
  "result" : [
    {
      "_key" : "9629",
      "_id" : "products/9629",
      "_rev" : "_U2FCCJu---",
      "hello1" : "world1"
    },
    {
      "_key" : "9633",
      "_id" : "products/9633",
      "_rev" : "_U2FCCJu--_",
      "hello2" : "world1"
    }
  ],
  "hasMore" : true,
  "id" : "9645",
  "count" : 5,
  "extra" : {
    "stats" : {
      "writesExecuted" : 0,
      "writesIgnored" : 0,
      "scannedFull" : 5,
      "scannedIndex" : 0,
      "filtered" : 0,
      "executionTime" : 0.00023603439331054688
    },
    "warnings" : [ ]
  },
  "cached" : false
}
```

Example:

Missing identifier

```
shell> curl -X PUT --dump - http://localhost:8529/_api/cursor

HTTP/1.1 400 Bad Request
content-type: application/json; charset=utf-8

{
  "error" : true,
  "errorMessage" : "expecting PUT /_api/cursor/<cursor-id>",
  "code" : 400,
  "errorNum" : 400
}
```

Example:

Unknown identifier

```
shell> curl -X PUT --dump - http://localhost:8529/_api/cursor/123123

HTTP/1.1 404 Not Found
content-type: application/json; charset=utf-8

{
  "error" : true,
  "errorMessage" : "cursor not found",
  "code" : 404,
  "errorNum" : 1600
}
```

Return Codes

- *200*: The server will respond with *HTTP 200* in case of success.
- *400*: If the cursor identifier is omitted, the server will respond with *HTTP 404*.
- *404*: If no cursor with the specified identifier can be found, the server will respond with *HTTP 404*.

Examples

Valid request for next batch

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/cursor <<EOF
{
  "query" : "FOR p IN products LIMIT 5 RETURN p",
  "count" : true,
  "batchSize" : 2
}
EOF

shell> curl -X PUT --dump - http://localhost:8529/_api/cursor/9645

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Missing identifier

```
shell> curl -X PUT --dump - http://localhost:8529/_api/cursor

HTTP/1.1 400 Bad Request
content-type: application/json; charset=utf-8
```

show response body

Unknown identifier

```
shell> curl -X PUT --dump - http://localhost:8529/_api/cursor/123123

HTTP/1.1 404 Not Found
content-type: application/json; charset=utf-8
```

show response body

Delete cursor

dispose an existing cursor

`DELETE /_api/cursor/{cursor-identifier}`

Path Parameters

- *cursor-identifier* (required): The id of the cursor

Deletes the cursor and frees the resources associated with it.

The cursor will automatically be destroyed on the server when the client has retrieved all documents from it. The client can also explicitly destroy the cursor at any earlier time using an HTTP DELETE request. The cursor id must be included as part of the URL.

Note: the server will also destroy abandoned cursors automatically after a certain server-controlled timeout to avoid resource leakage.

Example:

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/cursor <<EOF
{
  "query" : "FOR p IN products LIMIT 5 RETURN p",
  "count" : true,
  "batchSize" : 2
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8

{
  "error" : false,
  "code" : 201,
  "result" : [
    {
      "_key" : "9588",
      "_id" : "products/9588",
      "_rev" : "_U2FCBsu--A",
      "hello4" : "world1"
    },
    {
      "_key" : "9591",
      "_id" : "products/9591",
      "_rev" : "_U2FCBsy---",
      "hello5" : "world1"
    }
  ],
  "hasMore" : true,
  "id" : "9594",
  "count" : 5,
  "extra" : {
    "stats" : {
      "writesExecuted" : 0,
      "writesIgnored" : 0,
      "scannedFull" : 5,
      "scannedIndex" : 0,
      "filtered" : 0,
      "executionTime" : 0.0002009868621826172
    },
    "warnings" : [ ]
  },
  "cached" : false
}
shell> curl -X DELETE --dump - http://localhost:8529/_api/cursor/9594
```

Return Codes

- *202*: is returned if the server is aware of the cursor.
- *404*: is returned if the server is not aware of the cursor. It is also returned if a cursor is used after it has been destroyed.

Examples

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/cursor <<EOF
{
  "query" : "FOR p IN products LIMIT 5 RETURN p",
  "count" : true,
  "batchSize" : 2
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8
```

show response body

HTTP Interface for AQL Queries

Explaining and parsing queries

ArangoDB has an HTTP interface to syntactically validate AQL queries. Furthermore, it offers an HTTP interface to retrieve the execution plan for any valid AQL query.

Both functionalities do not actually execute the supplied AQL query, but only inspect it and return meta information about it.

Explain an AQL query

explain an AQL query and return information about it

`POST /_api/explain` A JSON object describing the query and query parameters.

A json post document with these Properties is required:

- **query**: the query which you want explained; If the query references any bind variables, these must also be passed in the attribute *bindVars*. Additional options for the query can be passed in the *options* attribute.
- **options**:
  - **optimizer.rules** (string): an array of to-be-included or to-be-excluded optimizer rules can be put into this attribute, telling the optimizer to include or exclude specific rules. To disable a rule, prefix its name with a `-`, to enable a rule, prefix it with a `+`. There is also a pseudo-rule `all`, which will match all optimizer rules.
  - **maxNumberOfPlans**: an optional maximum number of plans that the optimizer is allowed to generate. Setting this attribute to a low value allows to put a cap on the amount of work the optimizer does.
  - **allPlans**: if set to *true*, all possible execution plans will be returned. The default is *false*, meaning only the optimal plan will be returned.
- **bindVars** (object): key/value pairs representing the bind parameters.

To explain how an AQL query would be executed on the server, the query string can be sent to the server via an HTTP POST request. The server will then validate the query and create an execution plan for it. The execution plan will be returned, but the query will not be executed.

The execution plan that is returned by the server can be used to estimate the probable performance of the query. Though the actual performance will depend on many different factors, the execution plan normally can provide some rough estimates on the amount of work the server needs to do in order to actually run the query.

By default, the explain operation will return the optimal plan as chosen by the query optimizer The optimal plan is the plan with the lowest total estimated cost. The plan will be returned in the attribute *plan* of the response object. If the option *allPlans* is specified in the request, the result will contain all plans created by the optimizer. The plans will then be returned in the attribute *plans*.

The result will also contain an attribute *warnings*, which is an array of warnings that occurred during optimization or execution plan creation. Additionally, a *stats* attribute is contained in the result with some optimizer statistics. If *allPlans* is set to *false*, the result will contain an attribute *cacheable* that states whether the query results can be cached on the server if the query result cache were used. The *cacheable* attribute is not present when *allPlans* is set to *true*.

Each plan in the result is a JSON object with the following attributes:

- *nodes*: the array of execution nodes of the plan. The array of available node types can be found here
- *estimatedCost*: the total estimated cost for the plan. If there are multiple plans, the optimizer will choose the plan with the lowest total cost.
- *collections*: an array of collections used in the query
- *rules*: an array of rules the optimizer applied. An overview of the available rules can be found here
- *variables*: array of variables used in the query (note: this may contain internal variables created by the optimizer)

Example:

Valid query

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/explain <<EOF
{
  "query" : "FOR p IN products RETURN p"
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "plan" : {
    "nodes" : [
      {
        "type" : "SingletonNode",
        "dependencies" : [ ],
        "id" : 1,
        "estimatedCost" : 1,
        "estimatedNrItems" : 1
      },
      {
        "type" : "EnumerateCollectionNode",
        "dependencies" : [
          1
        ],
        "id" : 2,
        "estimatedCost" : 11,
        "estimatedNrItems" : 10,
        "database" : "_system",
        "collection" : "products",
        "outVariable" : {
          "id" : 0,
          "name" : "p"
        },
        "random" : false,
        "satellite" : false
      },
      {
        "type" : "ReturnNode",
        "dependencies" : [
          2
        ],
        "id" : 3,
        "estimatedCost" : 21,
        "estimatedNrItems" : 10,
        "inVariable" : {
          "id" : 0,
          "name" : "p"
        }
      }
    ],
    "rules" : [ ],
    "collections" : [
      {
        "name" : "products",
        "type" : "read"
      }
    ],
    "variables" : [
      {
        "id" : 0,
        "name" : "p"
      }
    ],
    "estimatedCost" : 21,
    "estimatedNrItems" : 10
  },
  "warnings" : [ ],
  "stats" : {
    "rulesExecuted" : 26,
    "rulesSkipped" : 0,
    "plansCreated" : 1
  },
  "cacheable" : true,
```

```
  "error" : false,
  "code" : 200
}
```

Example:

A plan with some optimizer rules applied

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/explain <<EOF
{
  "query" : "FOR p IN products LET a = p.id FILTER a == 4 LET name = p.name SORT p.id LIMIT 1 RETURN name"
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "plan" : {
    "nodes" : [
      {
        "type" : "SingletonNode",
        "dependencies" : [ ],
        "id" : 1,
        "estimatedCost" : 1,
        "estimatedNrItems" : 1
      },
      {
        "type" : "IndexNode",
        "dependencies" : [
          1
        ],
        "id" : 11,
        "estimatedCost" : 1.5,
        "estimatedNrItems" : 1,
        "database" : "_system",
        "collection" : "products",
        "satellite" : false,
        "outVariable" : {
          "id" : 0,
          "name" : "p"
        },
        "indexes" : [
          {
            "id" : "10133",
            "type" : "skiplist",
            "fields" : [
              "id"
            ],
            "unique" : false,
            "sparse" : false
          }
        ],
        "condition" : {
          "type" : "n-ary or",
          "subNodes" : [
            {
              "type" : "n-ary and",
              "subNodes" : [
                {
                  "type" : "compare ==",
                  "subNodes" : [
                    {
                      "type" : "attribute access",
                      "name" : "id",
                      "subNodes" : [
                        {
                          "type" : "reference",
                          "name" : "p",
                          "id" : 0
                        }
                      ]
```

```
                    },
                    {
                      "type" : "value",
                      "value" : 4
                    }
                  ]
                }
              ]
            }
          ]
        },
        "reverse" : false
      },
      {
        "type" : "CalculationNode",
        "dependencies" : [
          11
        ],
        "id" : 4,
        "estimatedCost" : 2.5,
        "estimatedNrItems" : 1,
        "expression" : {
          "type" : "compare ==",
          "subNodes" : [
            {
              "type" : "attribute access",
              "name" : "id",
              "subNodes" : [
                {
                  "type" : "reference",
                  "name" : "p",
                  "id" : 0
                }
              ]
            },
            {
              "type" : "value",
              "value" : 4
            }
          ]
        },
        "outVariable" : {
          "id" : 4,
          "name" : "3"
        },
        "canThrow" : false,
        "expressionType" : "simple"
      },
      {
        "type" : "FilterNode",
        "dependencies" : [
          4
        ],
        "id" : 5,
        "estimatedCost" : 3.5,
        "estimatedNrItems" : 1,
        "inVariable" : {
          "id" : 4,
          "name" : "3"
        }
      },
      {
        "type" : "LimitNode",
        "dependencies" : [
          5
        ],
        "id" : 9,
        "estimatedCost" : 4.5,
        "estimatedNrItems" : 1,
        "offset" : 0,
        "limit" : 1,
        "fullCount" : false
      },
      {
        "type" : "CalculationNode",
        "dependencies" : [
```

```
          9
        ],
        "id" : 6,
        "estimatedCost" : 5.5,
        "estimatedNrItems" : 1,
        "expression" : {
          "type" : "attribute access",
          "name" : "name",
          "subNodes" : [
            {
              "type" : "reference",
              "name" : "p",
              "id" : 0
            }
          ]
        },
        "outVariable" : {
          "id" : 2,
          "name" : "name"
        },
        "canThrow" : false,
        "expressionType" : "attribute"
      },
      {
        "type" : "ReturnNode",
        "dependencies" : [
          6
        ],
        "id" : 10,
        "estimatedCost" : 6.5,
        "estimatedNrItems" : 1,
        "inVariable" : {
          "id" : 2,
          "name" : "name"
        }
      }
    ],
    "rules" : [
      "move-calculations-up",
      "remove-redundant-calculations",
      "remove-unnecessary-calculations",
      "move-calculations-up-2",
      "use-indexes",
      "use-index-for-sort",
      "remove-unnecessary-calculations-2",
      "move-calculations-down"
    ],
    "collections" : [
      {
        "name" : "products",
        "type" : "read"
      }
    ],
    "variables" : [
      {
        "id" : 6,
        "name" : "5"
      },
      {
        "id" : 4,
        "name" : "3"
      },
      {
        "id" : 2,
        "name" : "name"
      },
      {
        "id" : 1,
        "name" : "a"
      },
      {
        "id" : 0,
        "name" : "p"
      }
    ],
    "estimatedCost" : 6.5,
```

```
    "estimatedNrItems" : 1
  },
  "warnings" : [ ],
  "stats" : {
    "rulesExecuted" : 26,
    "rulesSkipped" : 0,
    "plansCreated" : 1
  },
  "cacheable" : true,
  "error" : false,
  "code" : 200
}
```

Example:

Using some options

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/explain <<EOF
{
  "query" : "FOR p IN products LET a = p.id FILTER a == 4 LET name = p.name SORT p.id LIMIT 1 RETURN name",
  "options" : {
    "maxNumberOfPlans" : 2,
    "allPlans" : true,
    "optimizer" : {
      "rules" : [
        "-all",
        "+use-index-for-sort",
        "+use-index-range"
      ]
    }
  }
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "plans" : [
    {
      "nodes" : [
        {
          "type" : "SingletonNode",
          "dependencies" : [ ],
          "id" : 1,
          "estimatedCost" : 1,
          "estimatedNrItems" : 1
        },
        {
          "type" : "IndexNode",
          "dependencies" : [
            1
          ],
          "id" : 11,
          "estimatedCost" : 11,
          "estimatedNrItems" : 10,
          "database" : "_system",
          "collection" : "products",
          "satellite" : false,
          "outVariable" : {
            "id" : 0,
            "name" : "p"
          },
          "indexes" : [
            {
              "id" : "10171",
              "type" : "skiplist",
              "fields" : [
                "id"
              ],
              "unique" : false,
              "sparse" : false
```

```
            }
          ],
          "condition" : {
          },
          "reverse" : false
        },
        {
          "type" : "CalculationNode",
          "dependencies" : [
            11
          ],
          "id" : 3,
          "estimatedCost" : 21,
          "estimatedNrItems" : 10,
          "expression" : {
            "type" : "attribute access",
            "name" : "id",
            "subNodes" : [
              {
                "type" : "reference",
                "name" : "p",
                "id" : 0
              }
            ]
          },
          "outVariable" : {
            "id" : 1,
            "name" : "a"
          },
          "canThrow" : false,
          "expressionType" : "attribute"
        },
        {
          "type" : "CalculationNode",
          "dependencies" : [
            3
          ],
          "id" : 4,
          "estimatedCost" : 31,
          "estimatedNrItems" : 10,
          "expression" : {
            "type" : "compare ==",
            "subNodes" : [
              {
                "type" : "reference",
                "name" : "a",
                "id" : 1
              },
              {
                "type" : "value",
                "value" : 4
              }
            ]
          },
          "outVariable" : {
            "id" : 4,
            "name" : "3"
          },
          "canThrow" : false,
          "expressionType" : "simple"
        },
        {
          "type" : "FilterNode",
          "dependencies" : [
            4
          ],
          "id" : 5,
          "estimatedCost" : 41,
          "estimatedNrItems" : 10,
          "inVariable" : {
            "id" : 4,
            "name" : "3"
          }
        },
        {
          "type" : "CalculationNode",
```

```
          "dependencies" : [
            5
          ],
          "id" : 6,
          "estimatedCost" : 51,
          "estimatedNrItems" : 10,
          "expression" : {
            "type" : "attribute access",
            "name" : "name",
            "subNodes" : [
              {
                "type" : "reference",
                "name" : "p",
                "id" : 0
              }
            ]
          },
          "outVariable" : {
            "id" : 2,
            "name" : "name"
          },
          "canThrow" : false,
          "expressionType" : "attribute"
        },
        {
          "type" : "CalculationNode",
          "dependencies" : [
            6
          ],
          "id" : 7,
          "estimatedCost" : 61,
          "estimatedNrItems" : 10,
          "expression" : {
            "type" : "attribute access",
            "name" : "id",
            "subNodes" : [
              {
                "type" : "reference",
                "name" : "p",
                "id" : 0
              }
            ]
          },
          "outVariable" : {
            "id" : 6,
            "name" : "5"
          },
          "canThrow" : false,
          "expressionType" : "attribute"
        },
        {
          "type" : "LimitNode",
          "dependencies" : [
            7
          ],
          "id" : 9,
          "estimatedCost" : 62,
          "estimatedNrItems" : 1,
          "offset" : 0,
          "limit" : 1,
          "fullCount" : false
        },
        {
          "type" : "ReturnNode",
          "dependencies" : [
            9
          ],
          "id" : 10,
          "estimatedCost" : 63,
          "estimatedNrItems" : 1,
          "inVariable" : {
            "id" : 2,
            "name" : "name"
          }
        }
      ],
```

```
      "rules" : [
        "use-index-for-sort"
      ],
      "collections" : [
        {
          "name" : "products",
          "type" : "read"
        }
      ],
      "variables" : [
        {
          "id" : 6,
          "name" : "5"
        },
        {
          "id" : 4,
          "name" : "3"
        },
        {
          "id" : 2,
          "name" : "name"
        },
        {
          "id" : 1,
          "name" : "a"
        },
        {
          "id" : 0,
          "name" : "p"
        }
      ],
      "estimatedCost" : 63,
      "estimatedNrItems" : 1
    }
  ],
  "warnings" : [ ],
  "stats" : {
    "rulesExecuted" : 1,
    "rulesSkipped" : 25,
    "plansCreated" : 1
  },
  "error" : false,
  "code" : 200
}
```

Example:

Returning all plans

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/explain <<EOF
{
  "query" : "FOR p IN products FILTER p.id == 25 RETURN p",
  "options" : {
    "allPlans" : true
  }
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "plans" : [
    {
      "nodes" : [
        {
          "type" : "SingletonNode",
          "dependencies" : [ ],
          "id" : 1,
          "estimatedCost" : 1,
          "estimatedNrItems" : 1
        },
```

```
      {
        "type" : "IndexNode",
        "dependencies" : [
          1
        ],
        "id" : 6,
        "estimatedCost" : 1.99,
        "estimatedNrItems" : 1,
        "database" : "_system",
        "collection" : "products",
        "satellite" : false,
        "outVariable" : {
          "id" : 0,
          "name" : "p"
        },
        "indexes" : [
          {
            "id" : "10119",
            "type" : "hash",
            "fields" : [
              "id"
            ],
            "selectivityEstimate" : 1,
            "unique" : false,
            "sparse" : false
          }
        ],
        "condition" : {
          "type" : "n-ary or",
          "subNodes" : [
            {
              "type" : "n-ary and",
              "subNodes" : [
                {
                  "type" : "compare ==",
                  "subNodes" : [
                    {
                      "type" : "attribute access",
                      "name" : "id",
                      "subNodes" : [
                        {
                          "type" : "reference",
                          "name" : "p",
                          "id" : 0
                        }
                      ]
                    },
                    {
                      "type" : "value",
                      "value" : 25
                    }
                  ]
                }
              ]
            }
          ]
        },
        "reverse" : false
      },
      {
        "type" : "ReturnNode",
        "dependencies" : [
          6
        ],
        "id" : 5,
        "estimatedCost" : 2.99,
        "estimatedNrItems" : 1,
        "inVariable" : {
          "id" : 0,
          "name" : "p"
        }
      }
    ],
    "rules" : [
      "use-indexes",
      "remove-filter-covered-by-index",
```

```
        "remove-unnecessary-calculations-2"
      ],
      "collections" : [
        {
          "name" : "products",
          "type" : "read"
        }
      ],
      "variables" : [
        {
          "id" : 2,
          "name" : "1"
        },
        {
          "id" : 0,
          "name" : "p"
        }
      ],
      "estimatedCost" : 2.99,
      "estimatedNrItems" : 1
    }
  ],
  "warnings" : [ ],
  "stats" : {
    "rulesExecuted" : 26,
    "rulesSkipped" : 0,
    "plansCreated" : 1
  },
  "error" : false,
  "code" : 200
}
```

Example:

A query that produces a warning

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/explain <<EOF
{
  "query" : "FOR i IN 1..10 RETURN 1 / 0"
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "plan" : {
    "nodes" : [
      {
        "type" : "SingletonNode",
        "dependencies" : [ ],
        "id" : 1,
        "estimatedCost" : 1,
        "estimatedNrItems" : 1
      },
      {
        "type" : "CalculationNode",
        "dependencies" : [
          1
        ],
        "id" : 2,
        "estimatedCost" : 2,
        "estimatedNrItems" : 1,
        "expression" : {
          "type" : "range",
          "subNodes" : [
            {
              "type" : "value",
              "value" : 1
            },
            {
              "type" : "value",
```

```
              "value" : 10
            }
          ]
        },
        "outVariable" : {
          "id" : 2,
          "name" : "1"
        },
        "canThrow" : false,
        "expressionType" : "simple"
      },
      {
        "type" : "CalculationNode",
        "dependencies" : [
          2
        ],
        "id" : 4,
        "estimatedCost" : 3,
        "estimatedNrItems" : 1,
        "expression" : {
          "type" : "value",
          "value" : null
        },
        "outVariable" : {
          "id" : 4,
          "name" : "3"
        },
        "canThrow" : false,
        "expressionType" : "json"
      },
      {
        "type" : "EnumerateListNode",
        "dependencies" : [
          4
        ],
        "id" : 3,
        "estimatedCost" : 13,
        "estimatedNrItems" : 10,
        "inVariable" : {
          "id" : 2,
          "name" : "1"
        },
        "outVariable" : {
          "id" : 0,
          "name" : "i"
        }
      },
      {
        "type" : "ReturnNode",
        "dependencies" : [
          3
        ],
        "id" : 5,
        "estimatedCost" : 23,
        "estimatedNrItems" : 10,
        "inVariable" : {
          "id" : 4,
          "name" : "3"
        }
      }
    ],
    "rules" : [
      "move-calculations-up",
      "move-calculations-up-2"
    ],
    "collections" : [ ],
    "variables" : [
      {
        "id" : 4,
        "name" : "3"
      },
      {
        "id" : 2,
        "name" : "1"
      },
      {
```

```
        "id" : 0,
        "name" : "i"
      }
    ],
    "estimatedCost" : 23,
    "estimatedNrItems" : 10
  },
  "warnings" : [
    {
      "code" : 1562,
      "message" : "division by zero"
    }
  ],
  "stats" : {
    "rulesExecuted" : 26,
    "rulesSkipped" : 0,
    "plansCreated" : 1
  },
  "cacheable" : false,
  "error" : false,
  "code" : 200
}
```

Example:

Invalid query (missing bind parameter)

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/explain <<EOF
{
  "query" : "FOR p IN products FILTER p.id == @id LIMIT 2 RETURN p.n"
}
EOF

HTTP/1.1 400 Bad Request
content-type: application/json; charset=utf-8

{
  "error" : true,
  "code" : 400,
  "errorNum" : 1551,
  "errorMessage" : "no value specified for declared bind parameter 'id' (while parsing)"
}
```

Example:

The data returned in the **plan** attribute of the result contains one element per AQL top-level statement (i.e. `FOR`, `RETURN`, `FILTER` etc.). If the query optimizer removed some unnecessary statements, the result might also contain less elements than there were top-level statements in the AQL query.

The following example shows a query with a non-sensible filter condition that the optimizer has removed so that there are less top-level statements.

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/explain <<EOF
{ "query" : "FOR i IN [ 1, 2, 3 ] FILTER 1 == 2 RETURN i" }
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "plan" : {
    "nodes" : [
      {
        "type" : "SingletonNode",
        "dependencies" : [ ],
        "id" : 1,
        "estimatedCost" : 1,
```

```
        "estimatedNrItems" : 1
      },
      {
        "type" : "CalculationNode",
        "dependencies" : [
          1
        ],
        "id" : 2,
        "estimatedCost" : 2,
        "estimatedNrItems" : 1,
        "expression" : {
          "type" : "array",
          "subNodes" : [
            {
              "type" : "value",
              "value" : 1
            },
            {
              "type" : "value",
              "value" : 2
            },
            {
              "type" : "value",
              "value" : 3
            }
          ]
        },
        "outVariable" : {
          "id" : 2,
          "name" : "1"
        },
        "canThrow" : false,
        "expressionType" : "json"
      },
      {
        "type" : "NoResultsNode",
        "dependencies" : [
          2
        ],
        "id" : 7,
        "estimatedCost" : 0.5,
        "estimatedNrItems" : 0
      },
      {
        "type" : "EnumerateListNode",
        "dependencies" : [
          7
        ],
        "id" : 3,
        "estimatedCost" : 0.5,
        "estimatedNrItems" : 0,
        "inVariable" : {
          "id" : 2,
          "name" : "1"
        },
        "outVariable" : {
          "id" : 0,
          "name" : "i"
        }
      },
      {
        "type" : "ReturnNode",
        "dependencies" : [
          3
        ],
        "id" : 6,
        "estimatedCost" : 0.5,
        "estimatedNrItems" : 0,
        "inVariable" : {
          "id" : 0,
          "name" : "i"
        }
      }
    ],
    "rules" : [
      "move-calculations-up",
```

```
      "move-filters-up",
      "remove-unnecessary-filters",
      "remove-unnecessary-calculations"
    ],
    "collections" : [ ],
    "variables" : [
      {
        "id" : 4,
        "name" : "3"
      },
      {
        "id" : 2,
        "name" : "1"
      },
      {
        "id" : 0,
        "name" : "i"
      }
    ],
    "estimatedCost" : 0.5,
    "estimatedNrItems" : 0
  },
  "warnings" : [ ],
  "stats" : {
    "rulesExecuted" : 26,
    "rulesSkipped" : 0,
    "plansCreated" : 1
  },
  "cacheable" : true,
  "error" : false,
  "code" : 200
}
```

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/explain <<EOF
{
  "query" : "FOR p IN products RETURN p"
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

A plan with some optimizer rules applied

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/explain <<EOF
{
  "query" : "FOR p IN products LET a = p.id FILTER a == 4 LET name = p.name SORT p.id
LIMIT 1 RETURN name"
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Using some options

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/explain <<EOF
{
  "query" : "FOR p IN products LET a = p.id FILTER a == 4 LET name = p.name SORT p.id
LIMIT 1 RETURN name",
  "options" : {
    "maxNumberOfPlans" : 2,
    "allPlans" : true,
    "optimizer" : {
      "rules" : [
        "-all",
        "+use-index-for-sort",
        "+use-index-range"
      ]
    }
  }
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Returning all plans

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/explain <<EOF
{
  "query" : "FOR p IN products FILTER p.id == 25 RETURN p",
  "options" : {
    "allPlans" : true
  }
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

A query that produces a warning

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/explain <<EOF
{
  "query" : "FOR i IN 1..10 RETURN 1 / 0"
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Invalid query (missing bind parameter)

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/explain <<EOF
{
  "query" : "FOR p IN products FILTER p.id == @id LIMIT 2 RETURN p.n"
}
EOF

HTTP/1.1 400 Bad Request
content-type: application/json; charset=utf-8
```

show response body

The data returned in the **plan** attribute of the result contains one element per AQL top-level statement (i.e. `FOR`, `RETURN`, `FILTER` etc.). If the query optimizer removed some unnecessary statements, the result might also contain less elements than there were top-level statements in the AQL query. The following example shows a query with a non-sensible filter condition that the optimizer has removed so that there are less top-level statements.

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/explain <<EOF
{ "query" : "FOR i IN [ 1, 2, 3 ] FILTER 1 == 2 RETURN i" }
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Parse an AQL query

parse an AQL query and return information about it

`POST /_api/query`

This endpoint is for query validation only. To actually query the database, see `/api/cursor`.

A json post document with these Properties is required:

- **query**: To validate a query string without executing it, the query string can be passed to the server via an HTTP POST request.

Example:

a Valid query

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/query <<EOF
{ "query" : "FOR p IN products FILTER p.name == @name LIMIT 2 RETURN p.n" }
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "error" : false,
  "code" : 200,
  "parsed" : true,
  "collections" : [
    "products"
  ],
  "bindVars" : [
    "name"
  ],
  "ast" : [
    {
```

```
      "type" : "root",
      "subNodes" : [
        {
          "type" : "for",
          "subNodes" : [
            {
              "type" : "variable",
              "name" : "p",
              "id" : 0
            },
            {
              "type" : "collection",
              "name" : "products"
            }
          ]
        },
        {
          "type" : "filter",
          "subNodes" : [
            {
              "type" : "compare ==",
              "subNodes" : [
                {
                  "type" : "attribute access",
                  "name" : "name",
                  "subNodes" : [
                    {
                      "type" : "reference",
                      "name" : "p",
                      "id" : 0
                    }
                  ]
                },
                {
                  "type" : "parameter",
                  "name" : "name"
                }
              ]
            }
          ]
        },
        {
          "type" : "limit",
          "subNodes" : [
            {
              "type" : "value",
              "value" : 0
            },
            {
              "type" : "value",
              "value" : 2
            }
          ]
        },
        {
          "type" : "return",
          "subNodes" : [
            {
              "type" : "attribute access",
              "name" : "n",
              "subNodes" : [
                {
                  "type" : "reference",
                  "name" : "p",
                  "id" : 0
                }
              ]
            }
          ]
        }
      ]
    }
  ]
}
```

Example:

an Invalid query

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/query <<EOF
{ "query" : "FOR p IN products FILTER p.name = @name LIMIT 2 RETURN p.n" }
EOF

HTTP/1.1 400 Bad Request
content-type: application/json; charset=utf-8

{
  "error" : true,
  "errorMessage" : "syntax error, unexpected assignment near '= @name LIMIT 2 RETURN p.n' at position 1:33",
  "code" : 400,
  "errorNum" : 1501
}
```

Return Codes

- *200*: If the query is valid, the server will respond with *HTTP 200* and return the names of the bind parameters it found in the query (if any) in the *bindVars* attribute of the response. It will also return an array of the collections used in the query in the *collections* attribute. If a query can be parsed successfully, the *ast* attribute of the returned JSON will contain the abstract syntax tree representation of the query. The format of the *ast* is subject to change in future versions of ArangoDB, but it can be used to inspect how ArangoDB interprets a given query. Note that the abstract syntax tree will be returned without any optimizations applied to it.
- *400*: The server will respond with *HTTP 400* in case of a malformed request, or if the query contains a parse error. The body of the response will contain the error details embedded in a JSON object.

Examples

a Valid query

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/query <<EOF
{ "query" : "FOR p IN products FILTER p.name == @name LIMIT 2 RETURN p.n" }
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body
an Invalid query

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/query <<EOF
{ "query" : "FOR p IN products FILTER p.name = @name LIMIT 2 RETURN p.n" }
EOF

HTTP/1.1 400 Bad Request
content-type: application/json; charset=utf-8
```

show response body

Query tracking

ArangoDB has an HTTP interface for retrieving the lists of currently executing AQL queries and the list of slow AQL queries. In order to make meaningful use of these APIs, query tracking needs to be enabled in the database the HTTP request is executed for.

Returns the properties for the AQL query tracking

returns the configuration for the AQL query tracking

`GET /_api/query/properties`

Returns the current query tracking configuration. The configuration is a JSON object with the following properties:

- *enabled*: if set to *true*, then queries will be tracked. If set to *false*, neither queries nor slow queries will be tracked.
- *trackSlowQueries*: if set to *true*, then slow queries will be tracked in the list of slow queries if their runtime exceeds the value set in *slowQueryThreshold*. In order for slow queries to be tracked, the *enabled* property must also be set to *true*.
- *maxSlowQueries*: the maximum number of slow queries to keep in the list of slow queries. If the list of slow queries is full, the oldest entry in it will be discarded when additional slow queries occur.
- *slowQueryThreshold*: the threshold value for treating a query as slow. A query with a runtime greater or equal to this threshold value will be put into the list of slow queries when slow query tracking is enabled. The value for *slowQueryThreshold* is specified in seconds.
- *maxQueryStringLength*: the maximum query string length to keep in the list of queries. Query strings can have arbitrary lengths, and this property can be used to save memory in case very long query strings are used. The value is specified in bytes.

Return Codes

- *200*: Is returned if properties were retrieved successfully.
- *400*: The server will respond with *HTTP 400* in case of a malformed request,

Changes the properties for the AQL query tracking

changes the configuration for the AQL query tracking

`PUT /_api/query/properties`

A json post document with these Properties is required:

- **slowQueryThreshold**: The threshold value for treating a query as slow. A query with a runtime greater or equal to this threshold value will be put into the list of slow queries when slow query tracking is enabled. The value for *slowQueryThreshold* is specified in seconds.
- **enabled**: If set to *true*, then queries will be tracked. If set to *false*, neither queries nor slow queries will be tracked.
- **maxSlowQueries**: The maximum number of slow queries to keep in the list of slow queries. If the list of slow queries is full, the oldest entry in it will be discarded when additional slow queries occur.
- **trackSlowQueries**: If set to *true*, then slow queries will be tracked in the list of slow queries if their runtime exceeds the value set in *slowQueryThreshold*. In order for slow queries to be tracked, the *enabled* property must also be set to *true*.
- **maxQueryStringLength**: The maximum query string length to keep in the list of queries. Query strings can have arbitrary lengths, and this property can be used to save memory in case very long query strings are used. The value is specified in bytes.

The properties need to be passed in the attribute *properties* in the body of the HTTP request. *properties* needs to be a JSON object.

After the properties have been changed, the current set of properties will be returned in the HTTP response.

Return Codes

- *200*: Is returned if the properties were changed successfully.
- *400*: The server will respond with *HTTP 400* in case of a malformed request,

Returns the currently running AQL queries

returns a list of currently running AQL queries

`GET /_api/query/current`

Returns an array containing the AQL queries currently running in the selected database. Each query is a JSON object with the following attributes:

- *id*: the query's id
- *query*: the query string (potentially truncated)

- *started*: the date and time when the query was started
- *runTime*: the query's run time up to the point the list of queries was queried

Return Codes

- *200*: Is returned when the list of queries can be retrieved successfully.
- *400*: The server will respond with *HTTP 400* in case of a malformed request,

Returns the list of slow AQL queries

returns a list of slow running AQL queries

`GET /_api/query/slow`

Returns an array containing the last AQL queries that exceeded the slow query threshold in the selected database. The maximum amount of queries in the list can be controlled by setting the query tracking property `maxSlowQueries`. The threshold for treating a query as *slow* can be adjusted by setting the query tracking property `slowQueryThreshold`.

Each query is a JSON object with the following attributes:

- *id*: the query's id
- *query*: the query string (potentially truncated)
- *started*: the date and time when the query was started
- *runTime*: the query's run time up to the point the list of queries was queried

Return Codes

- *200*: Is returned when the list of queries can be retrieved successfully.
- *400*: The server will respond with *HTTP 400* in case of a malformed request,

Clears the list of slow AQL queries

clears the list of slow AQL queries

`DELETE /_api/query/slow`

Clears the list of slow AQL queries

Return Codes

- *200*: The server will respond with *HTTP 200* when the list of queries was cleared successfully.
- *400*: The server will respond with *HTTP 400* in case of a malformed request.

Killing queries

Running AQL queries can also be killed on the server. ArangoDB provides a kill facility via an HTTP interface. To kill a running query, its id (as returned for the query in the list of currently running queries) must be specified. The kill flag of the query will then be set, and the query will be aborted as soon as it reaches a cancelation point.

Kills a running AQL query

kills an AQL query

`DELETE /_api/query/{query-id}`

Path Parameters

- *query-id* (required): The id of the query.

Kills a running query. The query will be terminated at the next cancelation point.

Return Codes

- *200*: The server will respond with *HTTP 200* when the query was still running when the kill request was executed and the query's

kill flag was set.

- *400*: The server will respond with *HTTP 400* in case of a malformed request.
- *404*: The server will respond with *HTTP 404* when no query with the specified id was found.

HTTP Interface for the AQL query cache

This section describes the API methods for controlling the AQL query cache.

Clears any results in the AQL query cache

clears the AQL query cache

`DELETE /_api/query-cache`

clears the query cache

Return Codes

- *200*: The server will respond with *HTTP 200* when the cache was cleared successfully.
- *400*: The server will respond with *HTTP 400* in case of a malformed request.

Returns the global properties for the AQL query cache

returns the global configuration for the AQL query cache

`GET /_api/query-cache/properties`

Returns the global AQL query cache configuration. The configuration is a JSON object with the following properties:

- *mode*: the mode the AQL query cache operates in. The mode is one of the following values: *off*, *on* or *demand*.
- *maxResults*: the maximum number of query results that will be stored per database-specific cache.

Return Codes

- *200*: Is returned if the properties can be retrieved successfully.
- *400*: The server will respond with *HTTP 400* in case of a malformed request,

Globally adjusts the AQL query result cache properties

changes the configuration for the AQL query cache

`PUT /_api/query-cache/properties`

After the properties have been changed, the current set of properties will be returned in the HTTP response.

Note: changing the properties may invalidate all results in the cache. The global properties for AQL query cache. The properties need to be passed in the attribute *properties* in the body of the HTTP request. *properties* needs to be a JSON object with the following properties:

A json post document with these Properties is required:

- **mode**: the mode the AQL query cache should operate in. Possible values are *off*, *on* or *demand*.
- **maxResults**: the maximum number of query results that will be stored per database-specific cache.

Return Codes

- *200*: Is returned if the properties were changed successfully.
- *400*: The server will respond with *HTTP 400* in case of a malformed request,

HTTP Interface for AQL User Functions Management

AQL User Functions Management

This is an introduction to ArangoDB's HTTP interface for managing AQL user functions. AQL user functions are a means to extend the functionality of ArangoDB's query language (AQL) with user-defined JavaScript code.

For an overview of how AQL user functions and their implications, please refer to the Extending AQL chapter.

The HTTP interface provides an API for adding, deleting, and listing previously registered AQL user functions.

All user functions managed through this interface will be stored in the system collection *_aqlfunctions*. Documents in this collection should not be accessed directly, but only via the dedicated interfaces.

Create AQL user function

create a new AQL user function

`POST /_api/aqlfunction`

A json post document with these Properties is required:

- **isDeterministic**: an optional boolean value to indicate that the function results are fully deterministic (function return value solely depends on the input value and return value is the same for repeated calls with same input). The *isDeterministic* attribute is currently not used but may be used later for optimisations.
- **code**: a string representation of the function body.
- **name**: the fully qualified name of the user functions.

In case of success, the returned JSON object has the following properties:

- *error*: boolean flag to indicate that an error occurred (*false* in this case)
- *code*: the HTTP status code

The body of the response will contain a JSON object with additional error details. The object has the following attributes:

- *error*: boolean flag to indicate that an error occurred (*true* in this case)
- *code*: the HTTP status code
- *errorNum*: the server error number
- *errorMessage*: a descriptive error message

Example:

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/aqlfunction <<EOF
{
  "name" : "myfunctions::temperature::celsiustofahrenheit",
  "code" : "function (celsius) { return celsius * 1.8 + 32; }"
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8

{
  "error" : false,
  "code" : 201
}
```

Return Codes

- *200*: If the function already existed and was replaced by the call, the server will respond with *HTTP 200*.

- *201*: If the function can be registered by the server, the server will respond with *HTTP 201*.
- *400*: If the JSON representation is malformed or mandatory data is missing from the request, the server will respond with *HTTP 400*.

Examples

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/aqlfunction
<<EOF
{
  "name" : "myfunctions::temperature::celsiustofahrenheit",
  "code" : "function (celsius) { return celsius * 1.8 + 32; }"
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8

{
  "error" : false,
  "code" : 201
}
```

Remove existing AQL user function

remove an existing AQL user function

`DELETE /_api/aqlfunction/{name}`

Path Parameters

- *name* (required): the name of the AQL user function.

Query Parameters

- *group* (optional): If set to *true*, then the function name provided in *name* is treated as a namespace prefix, and all functions in the specified namespace will be deleted. If set to *false*, the function name provided in *name* must be fully qualified, including any namespaces.

Removes an existing AQL user function, identified by *name*.

In case of success, the returned JSON object has the following properties:

- *error*: boolean flag to indicate that an error occurred (*false* in this case)
- *code*: the HTTP status code

The body of the response will contain a JSON object with additional error details. The object has the following attributes:

- *error*: boolean flag to indicate that an error occurred (*true* in this case)
- *code*: the HTTP status code
- *errorNum*: the server error number
- *errorMessage*: a descriptive error message

Example:

deletes a function:

```
shell> curl -X DELETE --dump - http://localhost:8529/_api/aqlfunction/square::x::y

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "error" : false,
  "code" : 200
}
```

Example:

function not found:

```
shell> curl -X DELETE --dump - http://localhost:8529/_api/aqlfunction/myfunction::x::y

HTTP/1.1 404 Not Found
content-type: application/json; charset=utf-8

{
  "error" : true,
  "code" : 404,
  "errorNum" : 1582,
  "errorMessage" : "user function '%s()' not found"
}
```

Return Codes

- *200*: If the function can be removed by the server, the server will respond with *HTTP 200*.
- *400*: If the user function name is malformed, the server will respond with *HTTP 400*.
- *404*: If the specified user user function does not exist, the server will respond with *HTTP 404*.

Examples

deletes a function:

```
shell> curl -X DELETE --dump - http://localhost:8529/_api/aqlfunction/square::x::y

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "error" : false,
  "code" : 200
}
```

function not found:

```
shell> curl -X DELETE --dump - http://localhost:8529/_api/aqlfunction/myfunction::x::y

HTTP/1.1 404 Not Found
content-type: application/json; charset=utf-8
```

show response body

Return registered AQL user functions

gets all reqistered AQL user functions

`GET /_api/aqlfunction`

Query Parameters

- *namespace* (optional): Returns all registered AQL user functions from namespace *namespace*.

Returns all registered AQL user functions.

The call will return a JSON array with all user functions found. Each user function will at least have the following attributes:

- *name*: The fully qualified name of the user function
- *code*: A string representation of the function body

Example:

```
shell> curl --dump - http://localhost:8529/_api/aqlfunction

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

[
  {
    "name" : "myfunctions::temperature::celsiustofahrenheit",
    "code" : "function (celsius) { return celsius * 1.8 + 32; }"
  }
]
```

Return Codes

- *200*: if success *HTTP 200* is returned.

Examples

```
shell> curl --dump - http://localhost:8529/_api/aqlfunction

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

[
  {
    "name" : "myfunctions::temperature::celsiustofahrenheit",
    "code" : "function (celsius) { return celsius * 1.8 + 32; }"
  }
]
```

HTTP Interface for Simple Queries

Simple Queries

This is an introduction to ArangoDB's HTTP interface for simple queries.

Simple queries can be used if the query condition is straight forward simple, i.e., a document reference, all documents, a query-by-example, or a simple geo query. In a simple query you can specify exactly one collection and one condition. The result can then be sorted and can be split into pages.

Working with Simples Queries using HTTP

To limit the amount of results to be transferred in one batch, simple queries support a *batchSize* parameter that can optionally be used to tell the server to limit the number of results to be transferred in one batch to a certain value. If the query has more results than were transferred in one go, more results are waiting on the server so they can be fetched subsequently. If no value for the *batchSize* parameter is specified, the server will use a reasonable default value.

If the server has more documents than should be returned in a single batch, the server will set the *hasMore* attribute in the result. It will also return the id of the server-side cursor in the *id* attribute in the result. This id can be used with the cursor API to fetch any outstanding results from the server and dispose the server-side cursor afterwards.

Return all documents

returns all documents of a collection

`PUT /_api/simple/all`

Post Body

query (required): Contains the query.

free style json body

Returns all documents of a collections. The call expects a JSON object as body with the following attributes:

- *collection*: The name of the collection to query.
- *skip*: The number of documents to skip in the query (optional).
- *limit*: The maximal amount of documents to return. The *skip* is applied before the *limit* restriction. (optional)

Returns a cursor containing the result, see Http Cursor for details.

Example:

Limit the amount of documents using *limit*

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/all <<EOF
{ "collection": "products", "skip": 2, "limit" : 2 }
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8

{
  "result" : [
    {
      "_key" : "10629",
      "_id" : "products/10629",
      "_rev" : "_U2FCnmi---",
      "Hello4" : "World4"
    },
    {
      "_key" : "10623",
      "_id" : "products/10623",
      "_rev" : "_U2FCnme---",
      "Hello2" : "World2"
    }
  ],
  "hasMore" : false,
  "count" : 2,
  "cached" : false,
  "extra" : {
    "stats" : {
      "writesExecuted" : 0,
      "writesIgnored" : 0,
      "scannedFull" : 4,
      "scannedIndex" : 0,
      "filtered" : 0,
      "executionTime" : 0.00025010108947753906
    },
    "warnings" : [ ]
  },
  "error" : false,
  "code" : 201
}
```

Example:

Using a *batchSize* value

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/all <<EOF
{ "collection": "products", "batchSize" : 3 }
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8

{
  "error" : false,
  "code" : 201,
  "result" : [
    {
      "_key" : "10608",
      "_id" : "products/10608",
      "_rev" : "_U2FCndm---",
      "Hello4" : "World4"
    },
    {
      "_key" : "10611",
      "_id" : "products/10611",
      "_rev" : "_U2FCndm--_",
      "Hello5" : "World5"
    },
    {
      "_key" : "10605",
      "_id" : "products/10605",
      "_rev" : "_U2FCndi--_",
      "Hello3" : "World3"
    }
  ],
  "hasMore" : true,
  "id" : "10614",
  "count" : 5,
  "extra" : {
    "stats" : {
      "writesExecuted" : 0,
      "writesIgnored" : 0,
      "scannedFull" : 5,
      "scannedIndex" : 0,
      "filtered" : 0,
      "executionTime" : 0.00021195411682128906
    },
    "warnings" : [ ]
  },
  "cached" : false
}
```

Return Codes

- *201*: is returned if the query was executed successfully.
- *400*: is returned if the body does not contain a valid JSON representation of a query. The response body contains an error document in this case.
- *404*: is returned if the collection specified by *collection* is unknown. The response body contains an error document in this case.

Examples

Limit the amount of documents using *limit*

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/all <<EOF
{ "collection": "products", "skip": 2, "limit" : 2 }
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8
```

show response body

Using a *batchSize* value

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/all <<EOF
{ "collection": "products", "batchSize" : 3 }
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8
```

show response body

Simple query by-example

returns all documents of a collection matching a given example

`PUT /_api/simple/by-example`

A json post document with these Properties is required:

- **skip**: The number of documents to skip in the query (optional).
- **batchSize**: maximum number of result documents to be transferred from the server to the client in one roundtrip. If this attribute is not set, a server-controlled default value will be used. A *batchSize* value of *0* is disallowed.
- **limit**: The maximal amount of documents to return. The *skip* is applied before the *limit* restriction. (optional)
- **example**: The example document.
- **collection**: The name of the collection to query.

This will find all documents matching a given example.

Returns a cursor containing the result, see Http Cursor for details.

Example:

Matching an attribute

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/by-example <<EOF
{
  "collection" : "products",
  "example" : {
    "i" : 1
  }
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8

{
  "result" : [
    {
      "_key" : "10659",
      "_id" : "products/10659",
      "_rev" : "_U2FCn1W---",
      "a" : {
        "k" : 1,
        "j" : 1
      },
      "i" : 1
    },
    {
      "_key" : "10663",
      "_id" : "products/10663",
      "_rev" : "_U2FCn1a---",
      "a" : {
        "j" : 1
      },
      "i" : 1
    },
    {
      "_key" : "10666",
      "_id" : "products/10666",
      "_rev" : "_U2FCn1a--_",
      "i" : 1
    },
    {
      "_key" : "10669",
      "_id" : "products/10669",
      "_rev" : "_U2FCn1e---",
      "a" : {
        "k" : 2,
        "j" : 2
      },
      "i" : 1
    }
  ],
  "hasMore" : false,
  "count" : 4,
  "error" : false,
  "code" : 201
}
```

Example:

Matching an attribute which is a sub-document

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/by-example <<EOF
{
  "collection" : "products",
  "example" : {
    "a.j" : 1
  }
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8

{
  "result" : [
    {
      "_key" : "10677",
      "_id" : "products/10677",
      "_rev" : "_U2FCn8----",
      "a" : {
        "k" : 1,
        "j" : 1
      },
      "i" : 1
    },
    {
      "_key" : "10681",
      "_id" : "products/10681",
      "_rev" : "_U2FCn8---_",
      "a" : {
        "j" : 1
      },
      "i" : 1
    }
  ],
  "hasMore" : false,
  "count" : 2,
  "error" : false,
  "code" : 201
}
```

Example:

Matching an attribute within a sub-document

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/by-example <<EOF
{
  "collection" : "products",
  "example" : {
    "a" : {
      "j" : 1
    }
  }
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8

{
  "result" : [
    {
      "_key" : "10699",
      "_id" : "products/10699",
      "_rev" : "_U2FCoDW---",
      "a" : {
        "j" : 1
      },
      "i" : 1
    }
  ],
  "hasMore" : false,
  "count" : 1,
  "error" : false,
  "code" : 201
}
```

Return Codes

- *201*: is returned if the query was executed successfully.
- *400*: is returned if the body does not contain a valid JSON representation of a query. The response body contains an error document in this case.
- *404*: is returned if the collection specified by *collection* is unknown. The response body contains an error document in this case.

Examples

Matching an attribute

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/by-example <<EOF
{
  "collection" : "products",
  "example" : {
    "i" : 1
  }
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8
```

show response body

Matching an attribute which is a sub-document

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/by-
example <<EOF
{
  "collection" : "products",
  "example" : {
    "a.j" : 1
  }
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8
```

show response body

Matching an attribute within a sub-document

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/by-
example <<EOF
{
  "collection" : "products",
  "example" : {
    "a" : {
      "j" : 1
    }
  }
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8
```

show response body

Find documents matching an example

returns one document of a collection matching a given example

`PUT /_api/simple/first-example`

A json post document with these Properties is required:

- **example**: The example document.
- **collection**: The name of the collection to query.

This will return the first document matching a given example.

Returns a result containing the document or *HTTP 404* if no document matched the example.

If more than one document in the collection matches the specified example, only one of these documents will be returned, and it is undefined which of the matching documents is returned.

Example:

If a matching document was found

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/first-example <<EOF
{
  "collection" : "products",
  "example" : {
    "i" : 1
  }
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "document" : {
    "_key" : "10720",
    "_id" : "products/10720",
    "_rev" : "_U2FCoKu---",
    "i" : 1
  },
  "error" : false,
  "code" : 200
}
```

Example:

If no document was found

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/first-example <<EOF
{
  "collection" : "products",
  "example" : {
    "l" : 1
  }
}
EOF

HTTP/1.1 404 Not Found
content-type: application/json; charset=utf-8

{
  "error" : true,
  "code" : 404,
  "errorNum" : 404,
  "errorMessage" : "no match"
}
```

Return Codes

- *200*: is returned when the query was successfully executed.
- *400*: is returned if the body does not contain a valid JSON representation of a query. The response body contains an error document in this case.
- *404*: is returned if the collection specified by *collection* is unknown. The response body contains an error document in this case.

Examples

If a matching document was found

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/first-example <<EOF
{
  "collection" : "products",
  "example" : {
    "i" : 1
  }
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

If no document was found

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/first-example <<EOF
{
  "collection" : "products",
  "example" : {
    "l" : 1
  }
}
EOF

HTTP/1.1 404 Not Found
content-type: application/json; charset=utf-8
```

show response body

Find documents by their keys

fetches multiple documents by their keys

`PUT /_api/simple/lookup-by-keys`

A json post document with these Properties is required:

- **keys** (string): array with the _keys of documents to remove.
- **collection**: The name of the collection to look in for the documents

Looks up the documents in the specified collection using the array of keys provided. All documents for which a matching key was specified in the *keys* array and that exist in the collection will be returned. Keys for which no document can be found in the underlying collection are ignored, and no exception will be thrown for them.

The body of the response contains a JSON object with a *documents* attribute. The *documents* attribute is an array containing the matching documents. The order in which matching documents are present in the result array is unspecified.

Example:

Looking up existing documents

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/lookup-by-keys <<EOF
{
```

```
  "keys" : [
    "test0",
    "test1",
    "test2",
    "test3",
    "test4",
    "test5",
    "test6",
    "test7",
    "test8",
    "test9"
  ],
  "collection" : "test"
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "documents" : [
    {
      "_key" : "test0",
      "_id" : "test/test0",
      "_rev" : "_U2FCo2K---",
      "value" : 0
    },
    {
      "_key" : "test1",
      "_id" : "test/test1",
      "_rev" : "_U2FCo2K--_",
      "value" : 1
    },
    {
      "_key" : "test2",
      "_id" : "test/test2",
      "_rev" : "_U2FCo2O---",
      "value" : 2
    },
    {
      "_key" : "test3",
      "_id" : "test/test3",
      "_rev" : "_U2FCo2O--_",
      "value" : 3
    },
    {
      "_key" : "test4",
      "_id" : "test/test4",
      "_rev" : "_U2FCo2O--A",
      "value" : 4
    },
    {
      "_key" : "test5",
      "_id" : "test/test5",
      "_rev" : "_U2FCo2S---",
      "value" : 5
    },
    {
      "_key" : "test6",
      "_id" : "test/test6",
      "_rev" : "_U2FCo2S--_",
      "value" : 6
    },
    {
      "_key" : "test7",
      "_id" : "test/test7",
      "_rev" : "_U2FCo2S--A",
      "value" : 7
    },
    {
      "_key" : "test8",
      "_id" : "test/test8",
      "_rev" : "_U2FCo2W---",
      "value" : 8
    },
    {
```

```
      "_key" : "test9",
      "_id" : "test/test9",
      "_rev" : "_U2FCo2W--_",
      "value" : 9
    }
  ],
  "error" : false,
  "code" : 200
}
```

Example:

Looking up non-existing documents

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/lookup-by-keys <<EOF
{
  "keys" : [
    "foo",
    "bar",
    "baz"
  ],
  "collection" : "test"
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "documents" : [ ],
  "error" : false,
  "code" : 200
}
```

Return Codes

- *200*: is returned if the operation was carried out successfully.
- *404*: is returned if the collection was not found. The response body contains an error document in this case.
- *405*: is returned if the operation was called with a different HTTP METHOD than PUT.

Examples

Looking up existing documents

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/lookup-
by-keys <<EOF
{
  "keys" : [
    "test0",
    "test1",
    "test2",
    "test3",
    "test4",
    "test5",
    "test6",
    "test7",
    "test8",
    "test9"
  ],
  "collection" : "test"
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Looking up non-existing documents

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/lookup-
by-keys <<EOF
{
  "keys" : [
    "foo",
    "bar",
    "baz"
  ],
  "collection" : "test"
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Return a random document

returns a random document from a collection

`PUT /_api/simple/any`

Returns a random document from a collection. The call expects a JSON object as body with the following attributes:

A json post document with these Properties is required:

- **collection**: The identifier or name of the collection to query. Returns a JSON object with the document stored in the attribute *document* if the collection contains at least one document. If the collection is empty, the *document* attrbute contains null.

Example:

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/any <<EOF
{
  "collection" : "products"
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "document" : {
    "_key" : "10646",
    "_id" : "products/10646",
    "_rev" : "_U2FCnuW---",
    "Hello3" : "World3"
  },
  "error" : false,
  "code" : 200
}
```

Return Codes

- *200*: is returned if the query was executed successfully.
- *400*: is returned if the body does not contain a valid JSON representation of a query. The response body contains an error document in this case.
- *404*: is returned if the collection specified by *collection* is unknown. The response body contains an error document in this case.

Examples

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/any <<EOF
{
  "collection" : "products"
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Remove documents by their keys

removes multiple documents by their keys

`PUT /_api/simple/remove-by-keys`

A json post document with these Properties is required:

- **keys** (string): array with the _keys of documents to remove.
- **options**:
  - **returnOld**: if set to *true* and *silent* above is *false*, then the above information about the removed documents contains the complete removed documents.
  - **silent**: if set to *false*, then the result will contain an additional attribute *old* which contains an array with one entry for each removed document. By default, these entries will have the *_id*, *_key* and *_rev* attributes.
  - **waitForSync**: if set to true, then all removal operations will instantly be synchronized to disk. If this is not specified, then the collection's default sync behavior will be applied.
- **collection**: The name of the collection to look in for the documents to remove

Looks up the documents in the specified collection using the array of keys provided, and removes all documents from the collection whose keys are contained in the *keys* array. Keys for which no document can be found in the underlying collection are ignored, and no exception will be thrown for them.

The body of the response contains a JSON object with information how many documents were removed (and how many were not). The *removed* attribute will contain the number of actually removed documents. The *ignored* attribute will contain the number of keys in the request for which no matching document could be found.

Example:

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/remove-by-keys <<EOF
{
  "keys" : [
    "test0",
    "test1",
    "test2",
    "test3",
    "test4",
    "test5",
    "test6",
    "test7",
    "test8",
    "test9"
  ],
  "collection" : "test"
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "removed" : 10,
  "ignored" : 0,
  "error" : false,
  "code" : 200
}
```

Example:

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/remove-by-keys <<EOF
{
  "keys" : [
    "foo",
    "bar",
    "baz"
  ],
  "collection" : "test"
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "removed" : 0,
  "ignored" : 3,
  "error" : false,
  "code" : 200
}
```

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/remove-
by-keys <<EOF
{
  "keys" : [
    "test0",
    "test1",
    "test2",
    "test3",
    "test4",
    "test5",
    "test6",
    "test7",
    "test8",
    "test9"
  ],
  "collection" : "test"
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/remove-
by-keys <<EOF
{
  "keys" : [
    "foo",
    "bar",
    "baz"
  ],
  "collection" : "test"
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Remove documents by example

removes all documents of a collection that match an example

`PUT /_api/simple/remove-by-example`

A json post document with these Properties is required:

- **example**: An example document that all collection documents are compared against.
- **collection**: The name of the collection to remove from.
- **options**:
  - **limit**: an optional value that determines how many documents to delete at most. If *limit* is specified but is less than the number of documents in the collection, it is undefined which of the documents will be deleted.
  - **waitForSync**: if set to true, then all removal operations will instantly be synchronized to disk. If this is not specified, then the

collection's default sync behavior will be applied.

This will find all documents in the collection that match the specified example object.

Note: the *limit* attribute is not supported on sharded collections. Using it will result in an error.

Returns the number of documents that were deleted.

Example:

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/remove-by-example <<EOF
{
  "collection" : "products",
  "example" : {
    "a" : {
      "j" : 1
    }
  }
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "deleted" : 1,
  "error" : false,
  "code" : 200
}
```

Example:

Using Parameter: waitForSync and limit

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/remove-by-example <<EOF
{
  "collection" : "products",
  "example" : {
    "a" : {
      "j" : 1
    }
  },
  "waitForSync" : true,
  "limit" : 2
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "deleted" : 1,
  "error" : false,
  "code" : 200
}
```

Example:

Using Parameter: waitForSync and limit with new signature

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/remove-by-example <<EOF
{
  "collection" : "products",
  "example" : {
    "a" : {
      "j" : 1
    }
  },
  "options" : {
    "waitForSync" : true,
    "limit" : 2
  }
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "deleted" : 1,
  "error" : false,
  "code" : 200
}
```

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/remove-
by-example <<EOF
{
  "collection" : "products",
  "example" : {
    "a" : {
      "j" : 1
    }
  }
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Using Parameter: waitForSync and limit

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/remove-
by-example <<EOF
{
  "collection" : "products",
  "example" : {
    "a" : {
      "j" : 1
    }
  },
  "waitForSync" : true,
  "limit" : 2
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Using Parameter: waitForSync and limit with new signature

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/remove-
by-example <<EOF
{
  "collection" : "products",
  "example" : {
    "a" : {
      "j" : 1
    }
  },
  "options" : {
    "waitForSync" : true,
    "limit" : 2
  }
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Replace documents by example

replaces the body of all documents of a collection that match an example

`PUT /_api/simple/replace-by-example`

A json post document with these Properties is required:

- **options**:
  - **limit**: an optional value that determines how many documents to replace at most. If *limit* is specified but is less than the number of documents in the collection, it is undefined which of the documents will be replaced.
  - **waitForSync**: if set to true, then all removal operations will instantly be synchronized to disk. If this is not specified, then the collection's default sync behavior will be applied.
- **example**: An example document that all collection documents are compared against.
- **collection**: The name of the collection to replace within.
- **newValue**: The replacement document that will get inserted in place of the "old" documents.

This will find all documents in the collection that match the specified example object, and replace the entire document body with the new value specified. Note that document meta-attributes such as *_id*, *_key*, *_from*, *_to* etc. cannot be replaced.

Note: the *limit* attribute is not supported on sharded collections. Using it will result in an error.

Returns the number of documents that were replaced.

Example:

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/replace-by-example <<EOF
{
  "collection" : "products",
  "example" : {
    "a" : {
      "j" : 1
    }
  },
  "newValue" : {
    "foo" : "bar"
  },
  "limit" : 3
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "replaced" : 1,
  "error" : false,
  "code" : 200
}
```

Example:

Using new Signature for attributes WaitForSync and limit

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/replace-by-example <<EOF
{
  "collection" : "products",
  "example" : {
    "a" : {
      "j" : 1
    }
  },
  "newValue" : {
    "foo" : "bar"
  },
  "options" : {
    "limit" : 3,
    "waitForSync" : true
  }
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "replaced" : 1,
  "error" : false,
  "code" : 200
}
```

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/replace-
by-example <<EOF
{
  "collection" : "products",
  "example" : {
    "a" : {
      "j" : 1
    }
  },
  "newValue" : {
    "foo" : "bar"
  },
  "limit" : 3
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Using new Signature for attributes WaitForSync and limit

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/replace-
by-example <<EOF
{
  "collection" : "products",
  "example" : {
    "a" : {
      "j" : 1
    }
  },
  "newValue" : {
    "foo" : "bar"
  },
  "options" : {
    "limit" : 3,
    "waitForSync" : true
  }
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Update documents by example

partially updates the body of all documents of a collection that match an example

`PUT /_api/simple/update-by-example`

A json post document with these Properties is required:

- **options**:

  - **keepNull**: This parameter can be used to modify the behavior when handling *null* values. Normally, *null* values are stored in the database. By setting the *keepNull* parameter to *false*, this behavior can be changed so that all attributes in *data* with *null* values will be removed from the updated document.
  - **mergeObjects**: Controls whether objects (not arrays) will be merged if present in both the existing and the patch document. If set to false, the value in the patch document will overwrite the existing document's value. If set to true, objects will be merged. The default is true.
  - **limit**: an optional value that determines how many documents to update at most. If *limit* is specified but is less than the number of documents in the collection, it is undefined which of the documents will be updated.
  - **waitForSync**: if set to true, then all removal operations will instantly be synchronized to disk. If this is not specified, then the collection's default sync behavior will be applied.
- **example**: An example document that all collection documents are compared against.
- **collection**: The name of the collection to update within.
- **newValue**: A document containing all the attributes to update in the found documents.

This will find all documents in the collection that match the specified example object, and partially update the document body with the new value specified. Note that document meta-attributes such as *_id*, *_key*, *_from*, *_to* etc. cannot be replaced.

Note: the *limit* attribute is not supported on sharded collections. Using it will result in an error.

Returns the number of documents that were updated.

Example:

using old syntax for options

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/update-by-example <<EOF
{
  "collection" : "products",
  "example" : {
    "a" : {
      "j" : 1
    }
  },
  "newValue" : {
    "a" : {
      "j" : 22
    }
  },
  "limit" : 3
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "updated" : 1,
  "error" : false,
  "code" : 200
}
```

Example:

using new signature for options

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/update-by-example <<EOF
{
  "collection" : "products",
  "example" : {
    "a" : {
      "j" : 1
    }
  },
  "newValue" : {
    "a" : {
      "j" : 22
    }
  },
  "options" : {
    "limit" : 3,
    "waitForSync" : true
  }
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "updated" : 1,
  "error" : false,
  "code" : 200
}
```

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/update-
by-example <<EOF
{
  "collection" : "products",
  "example" : {
    "a" : {
      "j" : 1
    }
  },
  "newValue" : {
    "a" : {
      "j" : 22
    }
  },
  "limit" : 3
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

using new signature for options

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/update-
by-example <<EOF
{
  "collection" : "products",
  "example" : {
    "a" : {
      "j" : 1
    }
  },
  "newValue" : {
    "a" : {
      "j" : 22
    }
  },
  "options" : {
    "limit" : 3,
    "waitForSync" : true
  }
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Simple range query

returns all documents of a collection within a range

`PUT /_api/simple/range`

A json post document with these Properties is required:

- **right**: The upper bound.
- **attribute**: The attribute path to check.
- **collection**: The name of the collection to query.
- **limit**: The maximal amount of documents to return. The *skip* is applied before the *limit* restriction. (optional)
- **closed**: If *true*, use interval including *left* and *right*, otherwise exclude *right*, but include *left*.
- **skip**: The number of documents to skip in the query (optional).
- **left**: The lower bound.

This will find all documents within a given range. In order to execute a range query, a skip-list index on the queried attribute must be present.

Returns a cursor containing the result, see Http Cursor for details.

Note: the *range* simple query is **deprecated** as of ArangoDB 2.6. The function may be removed in future versions of ArangoDB. The preferred way for retrieving documents from a collection within a specific range is to use an AQL query as follows:

```
FOR doc IN @@collection
  FILTER doc.value >= @left && doc.value < @right
  LIMIT @skip, @limit
  RETURN doc`
```

Example:

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/range <<EOF
{
  "collection" : "products",
  "attribute" : "i",
  "left" : 2,
  "right" : 4
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8

{
  "result" : [
    {
      "_key" : "10918",
      "_id" : "products/10918",
      "_rev" : "_U2FCpl6--_",
      "i" : 2
    },
    {
      "_key" : "10921",
      "_id" : "products/10921",
      "_rev" : "_U2FCpm----",
      "i" : 3
    }
  ],
  "hasMore" : false,
  "count" : 2,
  "error" : false,
  "code" : 201
}
```

Return Codes

- *201*: is returned if the query was executed successfully.
- *400*: is returned if the body does not contain a valid JSON representation of a query. The response body contains an error document in this case.
- *404*: is returned if the collection specified by *collection* is unknown or no suitable index for the range query is present. The response body contains an error document in this case.

Examples

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/range
<<EOF
{
  "collection" : "products",
  "attribute" : "i",
  "left" : 2,
  "right" : 4
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8
```

show response body

Returns documents near a coordinate

returns all documents of a collection near a given location

`PUT /_api/simple/near`

A json post document with these Properties is required:

- **distance**: If given, the attribute key used to return the distance to the given coordinate. (optional). If specified, distances are returned in meters.
- **skip**: The number of documents to skip in the query. (optional)
- **longitude**: The longitude of the coordinate.
- **limit**: The maximal amount of documents to return. The *skip* is applied before the *limit* restriction. The default is 100. (optional)
- **collection**: The name of the collection to query.
- **latitude**: The latitude of the coordinate.
- **geo**: If given, the identifier of the geo-index to use. (optional)

The default will find at most 100 documents near the given coordinate. The returned array is sorted according to the distance, with the nearest document being first in the return array. If there are near documents of equal distance, documents are chosen randomly from this set until the limit is reached.

In order to use the *near* operator, a geo index must be defined for the collection. This index also defines which attribute holds the coordinates for the document. If you have more than one geo-spatial index, you can use the *geo* field to select a particular index.

Returns a cursor containing the result, see Http Cursor for details.

Note: the *near* simple query is **deprecated** as of ArangoDB 2.6. This API may be removed in future versions of ArangoDB. The preferred way for retrieving documents from a collection using the near operator is to issue an AQL query using the *NEAR* function as follows:

```
FOR doc IN NEAR(@@collection, @latitude, @longitude, @limit)
  RETURN doc`
```

Example:

Without distance

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/near <<EOF
{
  "collection" : "products",
  "latitude" : 0,
  "longitude" : 0,
  "skip" : 1,
  "limit" : 2
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8

{
  "result" : [
    {
      "_key" : "10847",
      "_id" : "products/10847",
      "_rev" : "_U2FCpKK--B",
      "name" : "Name/0.002/",
      "loc" : [
        0.002,
        0
      ]
    },
    {
      "_key" : "10841",
      "_id" : "products/10841",
      "_rev" : "_U2FCpKK--_",
      "name" : "Name/-0.002/",
      "loc" : [
        -0.002,
        0
      ]
    }
  ],
  "hasMore" : false,
  "count" : 2,
  "error" : false,
  "code" : 201
}
```

Example:

With distance

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/near <<EOF
{
  "collection" : "products",
  "latitude" : 0,
  "longitude" : 0,
  "skip" : 1,
  "limit" : 3,
  "distance" : "distance"
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8

{
  "result" : [
    {
      "distance" : 222.38985328911744,
      "_id" : "products/10884",
      "_key" : "10884",
      "_rev" : "_U2FCpXW---",
      "loc" : [
        -0.002,
        0
      ],
      "name" : "Name/-0.002/"
    },
    {
      "distance" : 222.38985328911744,
      "_id" : "products/10890",
      "_key" : "10890",
      "_rev" : "_U2FCpXW--A",
      "loc" : [
        0.002,
        0
      ],
      "name" : "Name/0.002/"
    },
    {
      "distance" : 444.779706578235,
      "_id" : "products/10881",
      "_key" : "10881",
      "_rev" : "_U2FCpXS--_",
      "loc" : [
        -0.004,
        0
      ],
      "name" : "Name/-0.004/"
    }
  ],
  "hasMore" : false,
  "count" : 3,
  "error" : false,
  "code" : 201
}
```

Return Codes

- *201*: is returned if the query was executed successfully.
- *400*: is returned if the body does not contain a valid JSON representation of a query. The response body contains an error document in this case.
- *404*: is returned if the collection specified by *collection* is unknown. The response body contains an error document in this case.

Examples

Without distance

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/near
<<EOF
{
  "collection" : "products",
  "latitude" : 0,
  "longitude" : 0,
  "skip" : 1,
  "limit" : 2
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8
```

show response body

With distance

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/near
<<EOF
{
  "collection" : "products",
  "latitude" : 0,
  "longitude" : 0,
  "skip" : 1,
  "limit" : 3,
  "distance" : "distance"
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8
```

show response body

Find documents within a radius around a coordinate

returns all documents of a collection within a given radius

`PUT /_api/simple/within`

A json post document with these Properties is required:

- **distance**: If given, the attribute key used to return the distance to the given coordinate. (optional). If specified, distances are returned in meters.
- **skip**: The number of documents to skip in the query. (optional)
- **longitude**: The longitude of the coordinate.
- **radius**: The maximal radius (in meters).
- **collection**: The name of the collection to query.
- **latitude**: The latitude of the coordinate.
- **limit**: The maximal amount of documents to return. The *skip* is applied before the *limit* restriction. The default is 100. (optional)
- **geo**: If given, the identifier of the geo-index to use. (optional)

This will find all documents within a given radius around the coordinate (*latitude*, *longitude*). The returned list is sorted by distance.

In order to use the *within* operator, a geo index must be defined for the collection. This index also defines which attribute holds the coordinates for the document. If you have more than one geo-spatial index, you can use the *geo* field to select a particular index.

Returns a cursor containing the result, see Http Cursor for details.

Note: the *within* simple query is **deprecated** as of ArangoDB 2.6. This API may be removed in future versions of ArangoDB. The preferred way for retrieving documents from a collection using the near operator is to issue an AQL query using the *WITHIN* function as follows:

```
FOR doc IN WITHIN(@@collection, @latitude, @longitude, @radius, @distanceAttributeName)
  RETURN doc
```

Example:

Without distance

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/near <<EOF
{
  "collection" : "products",
  "latitude" : 0,
  "longitude" : 0,
  "skip" : 1,
  "limit" : 2,
  "radius" : 500
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8

{
  "result" : [
    {
      "_key" : "11164",
      "_id" : "products/11164",
      "_rev" : "_U2FCrGO--_",
      "name" : "Name/0.002/",
      "loc" : [
        0.002,
        0
      ]
    },
    {
      "_key" : "11158",
      "_id" : "products/11158",
      "_rev" : "_U2FCrGK--_",
      "name" : "Name/-0.002/",
      "loc" : [
        -0.002,
        0
      ]
    }
  ],
  "hasMore" : false,
  "count" : 2,
  "error" : false,
  "code" : 201
}
```

Example:

With distance

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/near <<EOF
{
  "collection" : "products",
  "latitude" : 0,
  "longitude" : 0,
  "skip" : 1,
  "limit" : 3,
  "distance" : "distance",
  "radius" : 300
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8

{
  "result" : [
    {
      "distance" : 222.38985328911744,
      "_id" : "products/11201",
      "_key" : "11201",
      "_rev" : "_U2FCrSu--A",
      "loc" : [
        -0.002,
        0
      ],
      "name" : "Name/-0.002/"
    },
    {
      "distance" : 222.38985328911744,
      "_id" : "products/11207",
      "_key" : "11207",
      "_rev" : "_U2FCrSy--_",
      "loc" : [
        0.002,
        0
      ],
      "name" : "Name/0.002/"
    },
    {
      "distance" : 444.779706578235,
      "_id" : "products/11198",
      "_key" : "11198",
      "_rev" : "_U2FCrSu--_",
      "loc" : [
        -0.004,
        0
      ],
      "name" : "Name/-0.004/"
    }
  ],
  "hasMore" : false,
  "count" : 3,
  "error" : false,
  "code" : 201
}
```

Return Codes

- *201*: is returned if the query was executed successfully.
- *400*: is returned if the body does not contain a valid JSON representation of a query. The response body contains an error document in this case.
- *404*: is returned if the collection specified by *collection* is unknown. The response body contains an error document in this case.

Examples

Without distance

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/near
<<EOF
{
  "collection" : "products",
  "latitude" : 0,
  "longitude" : 0,
  "skip" : 1,
  "limit" : 2,
  "radius" : 500
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8
```

show response body

With distance

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/near
<<EOF
{
  "collection" : "products",
  "latitude" : 0,
  "longitude" : 0,
  "skip" : 1,
  "limit" : 3,
  "distance" : "distance",
  "radius" : 300
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8
```

show response body

Within rectangle query

returns all documents of a collection within a rectangle

`PUT /_api/simple/within-rectangle`

A json post document with these Properties is required:

- **latitude1**: The latitude of the first rectangle coordinate.
- **skip**: The number of documents to skip in the query. (optional)
- **latitude2**: The latitude of the second rectangle coordinate.
- **longitude2**: The longitude of the second rectangle coordinate.
- **longitude1**: The longitude of the first rectangle coordinate.
- **limit**: The maximal amount of documents to return. The *skip* is applied before the *limit* restriction. The default is 100. (optional)
- **collection**: The name of the collection to query.
- **geo**: If given, the identifier of the geo-index to use. (optional)

This will find all documents within the specified rectangle (determined by the given coordinates (*latitude1*, *longitude1*, *latitude2*, *longitude2*).

In order to use the *within-rectangle* query, a geo index must be defined for the collection. This index also defines which attribute holds the coordinates for the document. If you have more than one geo-spatial index, you can use the *geo* field to select a particular index.

Returns a cursor containing the result, see Http Cursor for details.

Example:

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/within-rectangle <<EOF
{
  "collection" : "products",
  "latitude1" : 0,
  "longitude1" : 0,
  "latitude2" : 0.2,
  "longitude2" : 0.2,
  "skip" : 1,
  "limit" : 2
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8

{
  "result" : [
    {
      "_key" : "11259",
      "_id" : "products/11259",
      "_rev" : "_U2FCrgi--A",
      "name" : "Name/0.008/",
      "loc" : [
        0.008,
        0
      ]
    },
    {
      "_key" : "11256",
      "_id" : "products/11256",
      "_rev" : "_U2FCrgi--_",
      "name" : "Name/0.006/",
      "loc" : [
        0.006,
        0
      ]
    }
  ],
  "hasMore" : false,
  "count" : 2,
  "error" : false,
  "code" : 201
}
```

Return Codes

- *201*: is returned if the query was executed successfully.
- *400*: is returned if the body does not contain a valid JSON representation of a query. The response body contains an error document in this case.
- *404*: is returned if the collection specified by *collection* is unknown. The response body contains an error document in this case.

Examples

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/within-rectangle <<EOF
{
  "collection" : "products",
  "latitude1" : 0,
  "longitude1" : 0,
  "latitude2" : 0.2,
  "longitude2" : 0.2,
  "skip" : 1,
  "limit" : 2
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8
```

show response body

Fulltext index query

returns documents of a collection as a result of a fulltext query

`PUT /_api/simple/fulltext`

A json post document with these Properties is required:

- **index**: The identifier of the fulltext-index to use.
- **attribute**: The attribute that contains the texts.
- **collection**: The name of the collection to query.
- **limit**: The maximal amount of documents to return. The *skip* is applied before the *limit* restriction. (optional)
- **skip**: The number of documents to skip in the query (optional).
- **query**: The fulltext query. Please refer to Fulltext queries for details.

This will find all documents from the collection that match the fulltext query specified in *query*.

In order to use the *fulltext* operator, a fulltext index must be defined for the collection and the specified attribute.

Returns a cursor containing the result, see Http Cursor for details.

Note: the *fulltext* simple query is **deprecated** as of ArangoDB 2.6. This API may be removed in future versions of ArangoDB. The preferred way for retrieving documents from a collection using the near operator is to issue an AQL query using the *FULLTEXT* AQL function as follows:

```
FOR doc IN FULLTEXT(@@collection, @attributeName, @queryString, @limit)
  RETURN doc
```

Example:

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/fulltext <<EOF
{
  "collection" : "products",
  "attribute" : "text",
  "query" : "word"
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8

{
  "result" : [
    {
      "_key" : "10752",
      "_id" : "products/10752",
      "_rev" : "_U2FCola--_",
      "text" : "this text also has a word"
    },
    {
      "_key" : "10748",
      "_id" : "products/10748",
      "_rev" : "_U2FCola---",
      "text" : "this text contains word"
    }
  ],
  "hasMore" : false,
  "count" : 2,
  "error" : false,
  "code" : 201
}
```

Return Codes

- *201*: is returned if the query was executed successfully.
- *400*: is returned if the body does not contain a valid JSON representation of a query. The response body contains an error document in this case.
- *404*: is returned if the collection specified by *collection* is unknown. The response body contains an error document in this case.

Examples

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/fulltext
<<EOF
{
  "collection" : "products",
  "attribute" : "text",
  "query" : "word"
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8
```

show response body

HTTP Interface for Async Results Management

Request Execution

ArangoDB provides various methods of executing client requests. Clients can choose the appropriate method on a per-request level based on their throughput, control flow, and durability requirements.

Blocking execution

ArangoDB is a multi-threaded server, allowing the processing of multiple client requests at the same time. Communication handling and the actual work can be performed by multiple worker threads in parallel.

Though multiple clients can connect and send their requests in parallel to ArangoDB, clients may need to wait for their requests to be processed.

By default, the server will fully process an incoming request and then return the result to the client. The client must wait for the server's response before it can send additional requests over the connection. For clients that are single-threaded or not event-driven, waiting for the full server response may be non-optimal.

Furthermore, please note that even if the client closes the HTTP connection, the request running on the server will still continue until it is complete and only then notice that the client no longer listens. Thus closing the connection does not help to abort a long running query! See below under Async Execution and later Result Retrieval and HttpJobPutCancel for details.

Fire and Forget

To mitigate client blocking issues, ArangoDB since version 1.4. offers a generic mechanism for non-blocking requests: if clients add the HTTP header *x-arango-async: true* to their requests, ArangoDB will put the request into an in-memory task queue and return an HTTP 202 (accepted) response to the client instantly. The server will execute the tasks from the queue asynchronously, decoupling the client requests and the actual work.

This allows for much higher throughput than if clients would wait for the server's response. The downside is that the response that is sent to the client is always the same (a generic HTTP 202) and clients cannot make a decision based on the actual operation's result at this point. In fact, the operation might have not even been executed at the time the generic response has reached the client. Clients can thus not rely on their requests having been processed successfully.

The asynchronous task queue on the server is not persisted, meaning not-yet processed tasks from the queue will be lost in case of a crash. However, the client will not know whether they were processed or not.

Clients should thus not send the extra header when they have strict durability requirements or if they rely on result of the sent operation for further actions.

The maximum number of queued tasks is determined by the startup option *-scheduler.maximal-queue-size*. If more than this number of tasks are already queued, the server will reject the request with an HTTP 500 error.

Finally, please note that it is not possible to cancel such a fire and forget job, since you won't get any handle to identify it later on. If you need to cancel requests, use Async Execution and later Result Retrieval and HttpJobPutCancel below.

Async Execution and later Result Retrieval

By adding the HTTP header *x-arango-async: store* to a request, clients can instruct the ArangoDB server to execute the operation asynchronously as above, but also store the operation result in memory for a later retrieval. The server will return a job id in the HTTP response header *x-arango-async-id*. The client can use this id in conjunction with the HTTP API at */_api/job*, which is described in detail in this manual.

Clients can ask the ArangoDB server via the async jobs API which results are ready for retrieval, and which are not. Clients can also use the async jobs API to retrieve the original results of an already executed async job by passing it the originally returned job id. The server will then return the job result as if the job was executed normally. Furthermore, clients can cancel running async jobs by their job id, see HttpJobPutCancel.

ArangoDB will keep all results of jobs initiated with the *x-arango-async: store* header. Results are removed from the server only if a client explicitly asks the server for a specific result.

The async jobs API also provides methods for garbage collection that clients can use to get rid of "old" not fetched results. Clients should call this method periodically because ArangoDB does not artificially limit the number of not-yet-fetched results.

It is thus a client responsibility to store only as many results as needed and to fetch available results as soon as possible, or at least to clean up not fetched results from time to time.

The job queue and the results are kept in memory only on the server, so they will be lost in case of a crash.

Canceling asynchronous jobs

As mentioned above it is possible to cancel an asynchronously running job using its job ID. This is done with a PUT request as described in HttpJobPutCancel.

However, a few words of explanation about what happens behind the scenes are in order. Firstly, a running async query can internally be executed by C++ code or by JavaScript code. For example CRUD operations are executed directly in C++, whereas AQL queries and transactions are executed by JavaScript code. The job cancelation only works for JavaScript code, since the mechanism used is simply to trigger an uncatchable exception in the JavaScript thread, which will be caught on the C++ level, which in turn leads to the cancelation of the job. No result can be retrieved later, since all data about the request is discarded.

If you cancel a job running on a coordinator of a cluster (Sharding), then only the code running on the coordinator is stopped, there may remain tasks within the cluster which have already been distributed to the DBservers and it is currently not possible to cancel them as well.

Async Execution and Authentication

If a request requires authentication, the authentication procedure is run before queueing. The request will only be queued if it valid credentials and the authentication succeeds. If the request does not contain valid credentials, it will not be queued but rejected instantly in the same way as a "regular", non-queued request.

Managing Async Results via HTTP

Return result of an async job

fetches a job result and removes it from the queue

`PUT /_api/job/{job-id}`

Path Parameters

- *job-id* (required): The async job id.

Returns the result of an async job identified by job-id. If the async job result is present on the server, the result will be removed from the list of result. That means this method can be called for each job-id once. The method will return the original job result's headers and body, plus the additional HTTP header x-arango-async-job-id. If this header is present, then the job was found and the response contains the original job's result. If the header is not present, the job was not found and the response contains status information from the job manager.

Example:

Not providing a job-id:

```
shell> curl -X PUT --dump - http://localhost:8529/_api/job

HTTP/1.1 400 Bad Request
content-type: application/json; charset=utf-8

{
  "error" : true,
  "errorMessage" : "bad parameter",
  "code" : 400,
  "errorNum" : 400
}
```

Example:

Providing a job-id for a non-existing job:

```
shell> curl -X PUT --dump - http://localhost:8529/_api/job/notthere

HTTP/1.1 404 Not Found
content-type: application/json; charset=utf-8

{
  "error" : true,
  "errorMessage" : "not found",
  "code" : 404,
  "errorNum" : 404
}
```

Example:

Fetching the result of an HTTP GET job:

```
shell> curl -X PUT --header 'x-arango-async: store' --dump - http://localhost:8529/_api/version

HTTP/1.1 202 Accepted
content-type: text/plain; charset=utf-8
x-arango-async-id: 149253070035374

shell> curl -X PUT --dump - http://localhost:8529/_api/job/149253070035374

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
x-arango-async-id: 149253070035374

{
  "server" : "arango",
  "version" : "3.1.18",
  "license" : "community"
}
```

Example:

Fetching the result of an HTTP POST job that failed:

```
shell> curl -X PUT --header 'x-arango-async: store' --data-binary @- --dump - http://localhost:8529/_api/collection <
<EOF
{
  "name" : " this name is invalid "
}
EOF

HTTP/1.1 202 Accepted
content-type: text/plain; charset=utf-8
x-arango-async-id: 149253070035379

shell> curl -X PUT --dump - http://localhost:8529/_api/job/149253070035379

HTTP/1.1 400 Bad Request
content-type: application/json; charset=utf-8
x-arango-async-id: 149253070035379

{
  "error" : true,
  "code" : 400,
  "errorNum" : 400,
  "errorMessage" : "expected PUT /_api/collection/<collection-name>/<action>"
}
```

Return Codes

- *204*: is returned if the job requested via job-id is still in the queue of pending (or not yet finished) jobs. In this case, no x-arango-async-id HTTP header will be returned.
- *400*: is returned if no job-id was specified in the request. In this case, no x-arango-async-id HTTP header will be returned.
- *404*: is returned if the job was not found or already deleted or fetched from the job result list. In this case, no x-arango-async-id HTTP header will be returned.

Examples

Not providing a job-id:

```
shell> curl -X PUT --dump - http://localhost:8529/_api/job

HTTP/1.1 400 Bad Request
content-type: application/json; charset=utf-8
```

show response body
Providing a job-id for a non-existing job:

```
shell> curl -X PUT --dump - http://localhost:8529/_api/job/notthere

HTTP/1.1 404 Not Found
content-type: application/json; charset=utf-8
```

show response body
Fetching the result of an HTTP GET job:

```
shell> curl -X PUT --header 'x-arango-async: store' --dump -
http://localhost:8529/_api/version

HTTP/1.1 202 Accepted
content-type: text/plain; charset=utf-8
x-arango-async-id: 149253070035374

shell> curl -X PUT --dump - http://localhost:8529/_api/job/149253070035374

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
x-arango-async-id: 149253070035374
```

show response body

Fetching the result of an HTTP POST job that failed:

```
shell> curl -X PUT --header 'x-arango-async: store' --data-binary @- --dump -
http://localhost:8529/_api/collection <<EOF
{
  "name" : " this name is invalid "
}
EOF

HTTP/1.1 202 Accepted
content-type: text/plain; charset=utf-8
x-arango-async-id: 149253070035379

shell> curl -X PUT --dump - http://localhost:8529/_api/job/149253070035379

HTTP/1.1 400 Bad Request
content-type: application/json; charset=utf-8
x-arango-async-id: 149253070035379
```

show response body

Cancel async job

cancels an async job

`PUT /_api/job/{job-id}/cancel`

Path Parameters

- *job-id* (required): The async job id.

Cancels the currently running job identified by job-id. Note that it still might take some time to actually cancel the running async job.

Example:

```
shell> curl -X POST --header 'x-arango-async: store' --data-binary @- --dump - http://localhost:8529/_api/cursor <<EO
F
{
  "query" : "FOR i IN 1..10 FOR j IN 1..10 LET x = sleep(1.0) FILTER i == 5 && j == 5 RETURN 42"
}
EOF

HTTP/1.1 202 Accepted
content-type: text/plain; charset=utf-8
x-arango-async-id: 149253070035337

shell> curl --dump - http://localhost:8529/_api/job/pending

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

[
  "149253070035337"
]
shell> curl -X PUT --dump - http://localhost:8529/_api/job/149253070035337/cancel

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "result" : true
}
shell> curl --dump - http://localhost:8529/_api/job/pending

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

[
  "149253070035337"
]
```

Return Codes

- *200*: cancel has been initiated.
- *400*: is returned if no job-id was specified in the request. In this case, no x-arango-async-id HTTP header will be returned.
- *404*: is returned if the job was not found or already deleted or fetched from the job result list. In this case, no x-arango-async-id HTTP header will be returned.

Examples

```
shell> curl -X POST --header 'x-arango-async: store' --data-binary @- --dump -
http://localhost:8529/_api/cursor <<EOF
{
  "query" : "FOR i IN 1..10 FOR j IN 1..10 LET x = sleep(1.0) FILTER i == 5 && j == 5
RETURN 42"
}
EOF

HTTP/1.1 202 Accepted
content-type: text/plain; charset=utf-8
x-arango-async-id: 149253070035337

shell> curl --dump - http://localhost:8529/_api/job/pending

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

[
  "149253070035337"
]
shell> curl -X PUT --dump - http://localhost:8529/_api/job/149253070035337/cancel

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Deletes async job

deletes an async job result

`DELETE /_api/job/{type}`

Path Parameters

- *type* (required): The type of jobs to delete. type can be:
- *all*: Deletes all jobs results. Currently executing or queued async jobs will not be stopped by this call.
- *expired*: Deletes expired results. To determine the expiration status of a result, pass the stamp query parameter. stamp needs to be a UNIX timestamp, and all async job results created at a lower timestamp will be deleted.
- *an actual job-id*: In this case, the call will remove the result of the specified async job. If the job is currently executing or queued, it will not be aborted.

Query Parameters

- *stamp* (optional): A UNIX timestamp specifying the expiration threshold when type is expired.

Deletes either all job results, expired job results, or the result of a specific job. Clients can use this method to perform an eventual garbage collection of job results.

Example:

Deleting all jobs:

```
shell> curl -X PUT --header 'x-arango-async: store' --dump - http://localhost:8529/_api/version

HTTP/1.1 202 Accepted
content-type: text/plain; charset=utf-8
x-arango-async-id: 149253070035346

shell> curl -X DELETE --dump - http://localhost:8529/_api/job/all

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "result" : true
}
```

Example:

Deleting expired jobs:

```
shell> curl -X PUT --header 'x-arango-async: store' --dump - http://localhost:8529/_api/version

HTTP/1.1 202 Accepted
content-type: text/plain; charset=utf-8
x-arango-async-id: 149253070035351

shell> curl --dump - http://localhost:8529/_admin/time

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "time" : 1492530721.019226,
  "error" : false,
  "code" : 200
}
shell> curl -X DELETE --dump - http://localhost:8529/_api/job/expired?stamp=1492530721.019226

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "result" : true
}
shell> curl --dump - http://localhost:8529/_api/job/pending

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

[ ]
```

Example:

Deleting the result of a specific job:

```
shell> curl -X PUT --header 'x-arango-async: store' --dump - http://localhost:8529/_api/version

HTTP/1.1 202 Accepted
content-type: text/plain; charset=utf-8
x-arango-async-id: 149253070035360

shell> curl -X DELETE --dump - http://localhost:8529/_api/job/149253070035360

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "result" : true
}
```

Example:

Deleting the result of a non-existing job:

```
shell> curl -X DELETE --dump - http://localhost:8529/_api/job/AreYouThere

HTTP/1.1 404 Not Found
content-type: application/json; charset=utf-8

{
  "error" : true,
  "errorMessage" : "not found",
  "code" : 404,
  "errorNum" : 404
}
```

Return Codes

- *200*: is returned if the deletion operation was carried out successfully. This code will also be returned if no results were deleted.
- *400*: is returned if type is not specified or has an invalid value.
- *404*: is returned if type is a job-id but no async job with the specified id was found.

Examples

Deleting all jobs:

```
shell> curl -X PUT --header 'x-arango-async: store' --dump -
http://localhost:8529/_api/version

HTTP/1.1 202 Accepted
content-type: text/plain; charset=utf-8
x-arango-async-id: 149253070035346

shell> curl -X DELETE --dump - http://localhost:8529/_api/job/all

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "result" : true
}
```

Deleting expired jobs:

```
shell> curl -X PUT --header 'x-arango-async: store' --dump -
http://localhost:8529/_api/version

HTTP/1.1 202 Accepted
content-type: text/plain; charset=utf-8
x-arango-async-id: 149253070035351

shell> curl --dump - http://localhost:8529/_admin/time

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Deleting the result of a specific job:

```
shell> curl -X PUT --header 'x-arango-async: store' --dump -
http://localhost:8529/_api/version

HTTP/1.1 202 Accepted
content-type: text/plain; charset=utf-8
x-arango-async-id: 149253070035360

shell> curl -X DELETE --dump - http://localhost:8529/_api/job/149253070035360

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "result" : true
}
```

Deleting the result of a non-existing job:

```
shell> curl -X DELETE --dump - http://localhost:8529/_api/job/AreYouThere

HTTP/1.1 404 Not Found
content-type: application/json; charset=utf-8
```

show response body

Returns async job

Returns the status of a specific job

`GET /_api/job/{job-id}`

Path Parameters

- *job-id* (required): The async job id.

Returns the processing status of the specified job. The processing status can be determined by peeking into the HTTP response code of the response.

Example:

Querying the status of a done job:

```
shell> curl -X PUT --header 'x-arango-async: store' --dump - http://localhost:8529/_api/version

HTTP/1.1 202 Accepted
content-type: text/plain; charset=utf-8
x-arango-async-id: 149253070035401

shell> curl -X PUT --dump - http://localhost:8529/_api/job/149253070035401

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
x-arango-async-id: 149253070035401

{
  "server" : "arango",
  "version" : "3.1.18",
  "license" : "community"
}
```

Example:

Querying the status of a pending job: (we create a sleep job therefore...)

```
shell> curl --header 'x-arango-async: store' --dump - http://localhost:8529/_admin/sleep?duration=30

HTTP/1.1 202 Accepted
content-type: text/plain; charset=utf-8
x-arango-async-id: 149253070035406

shell> curl --dump - http://localhost:8529/_api/job/149253070035406

HTTP/1.1 204 No Content
content-type: text/plain; charset=utf-8
```

Return Codes

- *200*: is returned if the job requested via job-id has been executed and its result is ready to fetch.
- *204*: is returned if the job requested via job-id is still in the queue of pending (or not yet finished) jobs.
- *404*: is returned if the job was not found or already deleted or fetched from the job result list.

Examples

Querying the status of a done job:

```
shell> curl -X PUT --header 'x-arango-async: store' --dump -
http://localhost:8529/_api/version

HTTP/1.1 202 Accepted
content-type: text/plain; charset=utf-8
x-arango-async-id: 149253070035401

shell> curl -X PUT --dump - http://localhost:8529/_api/job/149253070035401

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
x-arango-async-id: 149253070035401
```

show response body

Querying the status of a pending job: (we create a sleep job therefore...)

```
shell> curl --header 'x-arango-async: store' --dump -
http://localhost:8529/_admin/sleep?duration=30

HTTP/1.1 202 Accepted
content-type: text/plain; charset=utf-8
x-arango-async-id: 149253070035406

shell> curl --dump - http://localhost:8529/_api/job/149253070035406

HTTP/1.1 204 No Content
content-type: text/plain; charset=utf-8
```

Returns list of async jobs

Returns the ids of job results with a specific status

`GET /_api/job/{type}`

Path Parameters

- *type* (required): The type of jobs to return. The type can be either done or pending. Setting the type to done will make the method return the ids of already completed async jobs for which results can be fetched. Setting the type to pending will return the ids of not yet finished async jobs.

Query Parameters

- *count* (optional): The maximum number of ids to return per call. If not specified, a server-defined maximum value will be used.

Returns the list of ids of async jobs with a specific status (either done or pending). The list can be used by the client to get an overview of the job system status and to retrieve completed job results later.

Example:

Fetching the list of done jobs:

```
shell> curl -X PUT --header 'x-arango-async: store' --dump - http://localhost:8529/_api/version

HTTP/1.1 202 Accepted
content-type: text/plain; charset=utf-8
x-arango-async-id: 149253070035384

shell> curl --dump - http://localhost:8529/_api/job/done

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

[
  "149253070035384"
]
```

Example:

Fetching the list of pending jobs:

```
shell> curl -X PUT --header 'x-arango-async: store' --dump - http://localhost:8529/_api/version

HTTP/1.1 202 Accepted
content-type: text/plain; charset=utf-8
x-arango-async-id: 149253070035389

shell> curl --dump - http://localhost:8529/_api/job/pending

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

[ ]
```

Example:

Querying the status of a pending job: (we create a sleep job therefore...)

```
shell> curl --header 'x-arango-async: store' --dump - http://localhost:8529/_admin/sleep?duration=30

HTTP/1.1 202 Accepted
content-type: text/plain; charset=utf-8
x-arango-async-id: 149253070035394

shell> curl --dump - http://localhost:8529/_api/job/pending

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

[
  "149253070035394"
]
shell> curl -X DELETE --dump - http://localhost:8529/_api/job/149253070035394

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "result" : true
}
```

Return Codes

- *200*: is returned if the list can be compiled successfully. Note: the list might be empty.
- *400*: is returned if type is not specified or has an invalid value.

Examples

Fetching the list of done jobs:

```
shell> curl -X PUT --header 'x-arango-async: store' --dump -
http://localhost:8529/_api/version

HTTP/1.1 202 Accepted
content-type: text/plain; charset=utf-8
x-arango-async-id: 149253070035384

shell> curl --dump - http://localhost:8529/_api/job/done

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

[
  "149253070035384"
]
```

Fetching the list of pending jobs:

```
shell> curl -X PUT --header 'x-arango-async: store' --dump -
http://localhost:8529/_api/version

HTTP/1.1 202 Accepted
content-type: text/plain; charset=utf-8
x-arango-async-id: 149253070035389

shell> curl --dump - http://localhost:8529/_api/job/pending

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

[ ]
```

Querying the status of a pending job: (we create a sleep job therefore...)

```
shell> curl --header 'x-arango-async: store' --dump -
http://localhost:8529/_admin/sleep?duration=30

HTTP/1.1 202 Accepted
content-type: text/plain; charset=utf-8
x-arango-async-id: 149253070035394

shell> curl --dump - http://localhost:8529/_api/job/pending

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

[
  "149253070035394"
]
shell> curl -X DELETE --dump - http://localhost:8529/_api/job/149253070035394

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "result" : true
}
```

HTTP Interface for Bulk Imports

ArangoDB provides an HTTP interface to import multiple documents at once into a collection. This is known as a bulk import.

The data uploaded must be provided in JSON format. There are two mechanisms to import the data:

- self-contained JSON documents: in this case, each document contains all attribute names and values. Attribute names may be completely different among the documents uploaded
- attribute names plus document data: in this case, the first array must contain the attribute names of the documents that follow. The following arrays containing only the attribute values. Attribute values will be mapped to the attribute names by positions.

The endpoint address is */_api/import* for both input mechanisms. Data must be sent to this URL using an HTTP POST request. The data to import must be contained in the body of the POST request.

The *collection* query parameter must be used to specify the target collection for the import. Importing data into a non-existing collection will produce an error.

The *waitForSync* query parameter can be set to *true* to make the import only return if all documents have been synced to disk.

The *complete* query parameter can be set to *true* to make the entire import fail if any of the uploaded documents is invalid and cannot be imported. In this case, no documents will be imported by the import run, even if a failure happens at the end of the import.

If *complete* has a value other than *true*, valid documents will be imported while invalid documents will be rejected, meaning only some of the uploaded documents might have been imported.

The *details* query parameter can be set to *true* to make the import API return details about documents that could not be imported. If *details* is *true*, then the result will also contain a *details* attribute which is an array of detailed error messages. If the *details* is set to *false* or omitted, no details will be returned.

imports document values

imports documents from JSON-encoded lists

`POST /_api/import#document`

Post Body

documents (required): The body must consist of JSON-encoded arrays of attribute values, with one line per document. The first row of the request must be a JSON-encoded array of attribute names. These attribute names are used for the data in the subsequent lines.

Query Parameters

- *collection* (required): The collection name.
- *fromPrefix* (optional): An optional prefix for the values in `_from` attributes. If specified, the value is automatically prepended to each `_from` input value. This allows specifying just the keys for `_from`.
- *toPrefix* (optional): An optional prefix for the values in `_to` attributes. If specified, the value is automatically prepended to each `_to` input value. This allows specifying just the keys for `_to`.
- *overwrite* (optional): If this parameter has a value of `true` or `yes`, then all data in the collection will be removed prior to the import. Note that any existing index definitions will be preseved.
- *waitForSync* (optional): Wait until documents have been synced to disk before returning.
- *onDuplicate* (optional): Controls what action is carried out in case of a unique key constraint violation. Possible values are:
- *error*: this will not import the current document because of the unique key constraint violation. This is the default setting.
- *update*: this will update an existing document in the database with the data specified in the request. Attributes of the existing document that are not present in the request will be preseved.
- *replace*: this will replace an existing document in the database with the data specified in the request.
- *ignore*: this will not update an existing document and simply ignore the error caused by the unique key constraint violation. Note that *update*, *replace* and *ignore* will only work when the import document in the request contains the *_key* attribute. *update* and *replace* may also fail because of secondary unique key constraint violations.
- *complete* (optional): If set to `true` or `yes`, it will make the whole import fail if any error occurs. Otherwise the import will continue even if some documents cannot be imported.
- *details* (optional): If set to `true` or `yes`, the result will include an attribute `details` with details about documents that could not

be imported.

free style json body

NOTE Swagger examples won't work due to the anchor.

Creates documents in the collection identified by `collection-name`. The first line of the request body must contain a JSON-encoded array of attribute names. All following lines in the request body must contain JSON-encoded arrays of attribute values. Each line is interpreted as a separate document, and the values specified will be mapped to the array of attribute names specified in the first header line.

The response is a JSON object with the following attributes:

- `created`: number of documents imported.
- `errors`: number of documents that were not imported due to an error.
- `empty`: number of empty lines found in the input (will only contain a value greater zero for types `documents` or `auto`).
- `updated`: number of updated/replaced documents (in case `onDuplicate` was set to either `update` or `replace`).
- `ignored`: number of failed but ignored insert operations (in case `onDuplicate` was set to `ignore`).
- `details`: if query parameter `details` is set to true, the result will contain a `details` attribute which is an array with more detailed information about which documents could not be inserted.

Note: this API is currently not supported on cluster coordinators.

Example:

Importing two documents, with attributes `_key`, `value1` and `value2` each. One line in the import data is empty

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/import?collection=products <<EOF
[ "_key", "value1", "value2" ]
[ "abc", 25, "test" ]

[ "foo", "bar", "baz" ]
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8

{
  "error" : false,
  "created" : 2,
  "errors" : 0,
  "empty" : 1,
  "updated" : 0,
  "ignored" : 0
}
```

Example:

Importing into an edge collection, with attributes `_from`, `_to` and `name`

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/import?collection=links <<EOF
[ "_from", "_to", "name" ]
[ "products/123","products/234", "some name" ]
[ "products/332", "products/abc", "other name" ]
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8

{
  "error" : false,
  "created" : 2,
  "errors" : 0,
  "empty" : 0,
  "updated" : 0,
  "ignored" : 0
}
```

Example:

Importing into an edge collection, omitting `_from` or `_to`

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/import?collection=links&details=true <<EOF
[ "name" ]
[ "some name" ]
[ "other name" ]
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8

{
  "error" : false,
  "created" : 0,
  "errors" : 2,
  "empty" : 0,
  "updated" : 0,
  "ignored" : 0,
  "details" : [
    "at position 1: missing '_from' or '_to' attribute, offending document: {\"name\":\"some name\"}",
    "at position 2: missing '_from' or '_to' attribute, offending document: {\"name\":\"other name\"}"
  ]
}
```

Example:

Violating a unique constraint, but allow partial imports

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/import?collection=products&details=true <<EO
F
[ "_key", "value1", "value2" ]
[ "abc", 25, "test" ]
["abc", "bar", "baz" ]
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8

{
  "error" : false,
  "created" : 1,
  "errors" : 1,
  "empty" : 0,
  "updated" : 0,
  "ignored" : 0,
  "details" : [
    "at position 1: creating document failed with error 'unique constraint violated', offending document: {\"_key\":\
"abc\",\"value1\":\"bar\",\"value2\":\"baz\"}"
  ]
}
```

Example:

Violating a unique constraint, not allowing partial imports

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/import?collection=products&complete=true <<E
OF
[ "_key", "value1", "value2" ]
[ "abc", 25, "test" ]
["abc", "bar", "baz" ]
EOF

HTTP/1.1 409 Conflict
content-type: application/json; charset=utf-8

{
  "error" : true,
  "errorMessage" : "cannot create document, unique constraint violated",
  "code" : 409,
  "errorNum" : 1210
}
```

Example:

Using a non-existing collection

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/import?collection=products <<EOF
[ "_key", "value1", "value2" ]
[ "abc", 25, "test" ]
["foo", "bar", "baz" ]
EOF

HTTP/1.1 404 Not Found
content-type: application/json; charset=utf-8

{
  "error" : true,
  "errorMessage" : "collection 'products' not found",
  "code" : 404,
  "errorNum" : 1203
}
```

Example:

Using a malformed body

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/import?collection=products <<EOF
{ "_key": "foo", "value1": "bar" }
EOF

HTTP/1.1 400 Bad Request
content-type: application/json; charset=utf-8

{
  "error" : true,
  "errorMessage" : "no JSON array found in second line",
  "code" : 400,
  "errorNum" : 400
}
```

Return Codes

- *201*: is returned if all documents could be imported successfully.
- *400*: is returned if `type` contains an invalid value, no `collection` is specified, the documents are incorrectly encoded, or the request is malformed.
- *404*: is returned if `collection` or the `_from` or `_to` attributes of an imported edge refer to an unknown collection.
- *409*: is returned if the import would trigger a unique key violation and `complete` is set to `true`.
- *500*: is returned if the server cannot auto-generate a document key (out of keys error) for a document with no user-defined key.
- *501*: The server will respond with *HTTP 501* if this API is called on a cluster coordinator.

Examples

Importing two documents, with attributes `_key`, `value1` and `value2` each. One line in the import data is empty

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/import?
collection=products <<EOF
[ "_key", "value1", "value2" ]
[ "abc", 25, "test" ]

[ "foo", "bar", "baz" ]
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8
```

show response body

Importing into an edge collection, with attributes `_from`, `_to` and `name`

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/import?
collection=links <<EOF
[ "_from", "_to", "name" ]
[ "products/123","products/234", "some name" ]
[ "products/332", "products/abc", "other name" ]
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8
```

show response body

Importing into an edge collection, omitting `_from` or `_to`

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/import?
collection=links&details=true <<EOF
[ "name" ]
[ "some name" ]
[ "other name" ]
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8
```

show response body

Violating a unique constraint, but allow partial imports

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/import?
collection=products&details=true <<EOF
[ "_key", "value1", "value2" ]
[ "abc", 25, "test" ]
["abc", "bar", "baz" ]
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8
```

show response body

Violating a unique constraint, not allowing partial imports

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/import?
collection=products&complete=true <<EOF
[ "_key", "value1", "value2" ]
[ "abc", 25, "test" ]
["abc", "bar", "baz" ]
EOF

HTTP/1.1 409 Conflict
content-type: application/json; charset=utf-8
```

show response body

Using a non-existing collection

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/import?
collection=products <<EOF
[ "_key", "value1", "value2" ]
[ "abc", 25, "test" ]
["foo", "bar", "baz" ]
EOF

HTTP/1.1 404 Not Found
content-type: application/json; charset=utf-8
```

show response body

Using a malformed body

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/import?
collection=products <<EOF
{ "_key": "foo", "value1": "bar" }
EOF

HTTP/1.1 400 Bad Request
content-type: application/json; charset=utf-8
```

show response body

imports documents from JSON

imports documents from JSON

`POST /_api/import#json`

Post Body

documents (required): The body must either be a JSON-encoded array of objects or a string with multiple JSON objects separated by newlines.

Query Parameters

- *type* (required): Determines how the body of the request will be interpreted. `type` can have the following values:
- `documents` : when this type is used, each line in the request body is expected to be an individual JSON-encoded document. Multiple JSON objects in the request body need to be separated by newlines.
- `list` : when this type is used, the request body must contain a single JSON-encoded array of individual objects to import.
- `auto` : if set, this will automatically determine the body type (either `documents` or `list`).
- *collection* (required): The collection name.
- *fromPrefix* (optional): An optional prefix for the values in `_from` attributes. If specified, the value is automatically prepended to each `_from` input value. This allows specifying just the keys for `_from` .
- *toPrefix* (optional): An optional prefix for the values in `_to` attributes. If specified, the value is automatically prepended to each `_to` input value. This allows specifying just the keys for `_to` .
- *overwrite* (optional): If this parameter has a value of `true` or `yes` , then all data in the collection will be removed prior to the import. Note that any existing index definitions will be preseved.
- *waitForSync* (optional): Wait until documents have been synced to disk before returning.
- *onDuplicate* (optional): Controls what action is carried out in case of a unique key constraint violation. Possible values are:
- *error*: this will not import the current document because of the unique key constraint violation. This is the default setting.
- *update*: this will update an existing document in the database with the data specified in the request. Attributes of the existing document that are not present in the request will be preseved.
- *replace*: this will replace an existing document in the database with the data specified in the request.
- *ignore*: this will not update an existing document and simply ignore the error caused by a unique key constraint violation. Note that that *update*, *replace* and *ignore* will only work when the import document in the request contains the *_key* attribute. *update* and *replace* may also fail because of secondary unique key constraint violations.
- *complete* (optional): If set to `true` or `yes` , it will make the whole import fail if any error occurs. Otherwise the import will continue even if some documents cannot be imported.
- *details* (optional): If set to `true` or `yes` , the result will include an attribute `details` with details about documents that could not be imported.

free style json body

NOTE Swagger examples won't work due to the anchor.

Creates documents in the collection identified by `collection-name` . The JSON representations of the documents must be passed as the body of the POST request. The request body can either consist of multiple lines, with each line being a single stand-alone JSON object, or a singe JSON array with sub-objects.

The response is a JSON object with the following attributes:

- `created` : number of documents imported.
- `errors` : number of documents that were not imported due to an error.
- `empty` : number of empty lines found in the input (will only contain a value greater zero for types `documents` or `auto`).
- `updated` : number of updated/replaced documents (in case `onDuplicate` was set to either `update` or `replace`).
- `ignored` : number of failed but ignored insert operations (in case `onDuplicate` was set to `ignore`).
- `details` : if query parameter `details` is set to true, the result will contain a `details` attribute which is an array with more detailed information about which documents could not be inserted.

Note: this API is currently not supported on cluster coordinators.

Example:

Importing documents with heterogenous attributes from a JSON array

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/import?collection=products&type=list <<EOF
[
  {
    "_key" : "abc",
    "value1" : 25,
    "value2" : "test",
    "allowed" : true
  },
  {
    "_key" : "foo",
    "name" : "baz"
  },
  {
    "name" : {
      "detailed" : "detailed name",
      "short" : "short name"
    }
  }
]
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8

{
  "error" : false,
  "created" : 3,
  "errors" : 0,
  "empty" : 0,
  "updated" : 0,
  "ignored" : 0
}
```

Example:

Importing documents from individual JSON lines

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/import?collection=products&type=documents <<
EOF
{ "_key": "abc", "value1": 25, "value2": "test","allowed": true }
{ "_key": "foo", "name": "baz" }

{ "name": { "detailed": "detailed name", "short": "short name" } }

EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8

{
  "error" : false,
  "created" : 3,
  "errors" : 0,
  "empty" : 1,
  "updated" : 0,
  "ignored" : 0
}
```

Example:

Using the auto type detection

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/import?collection=products&type=auto <<EOF
[
  {
    "_key" : "abc",
    "value1" : 25,
    "value2" : "test",
    "allowed" : true
  },
  {
    "_key" : "foo",
    "name" : "baz"
  },
  {
    "name" : {
      "detailed" : "detailed name",
      "short" : "short name"
    }
  }
]
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8

{
  "error" : false,
  "created" : 3,
  "errors" : 0,
  "empty" : 0,
  "updated" : 0,
  "ignored" : 0
}
```

Example:

Importing into an edge collection, with attributes `_from`, `_to` and `name`

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/import?collection=links&type=documents <<EOF
{ "_from": "products/123", "_to": "products/234" }
{"_from": "products/332", "_to": "products/abc",   "name": "other name" }
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8

{
  "error" : false,
  "created" : 2,
  "errors" : 0,
  "empty" : 0,
  "updated" : 0,
  "ignored" : 0
}
```

Example:

Importing into an edge collection, omitting `_from` or `_to`

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/import?collection=links&type=list&details=tr
ue <<EOF
[
  {
    "name" : "some name"
  }
]
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8

{
  "error" : false,
  "created" : 0,
  "errors" : 1,
  "empty" : 0,
  "updated" : 0,
  "ignored" : 0,
  "details" : [
    "at position 1: missing '_from' or '_to' attribute, offending document: {\"name\":\"some name\"}"
  ]
}
```

Example:

Violating a unique constraint, but allow partial imports

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/import?collection=products&type=documents&de
tails=true <<EOF
{ "_key": "abc", "value1": 25, "value2": "test" }
{ "_key": "abc", "value1": "bar", "value2": "baz" }
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8

{
  "error" : false,
  "created" : 1,
  "errors" : 1,
  "empty" : 0,
  "updated" : 0,
  "ignored" : 0,
  "details" : [
    "at position 1: creating document failed with error 'unique constraint violated', offending document: {\"_key\":\
"abc\",\"value1\":\"bar\",\"value2\":\"baz\"}"
  ]
}
```

Example:

Violating a unique constraint, not allowing partial imports

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/import?collection=products&type=documents&co
mplete=true <<EOF
{ "_key": "abc", "value1": 25, "value2": "test" }
{ "_key": "abc", "value1": "bar", "value2": "baz" }
EOF

HTTP/1.1 409 Conflict
content-type: application/json; charset=utf-8

{
  "error" : true,
  "errorMessage" : "cannot create document, unique constraint violated",
  "code" : 409,
  "errorNum" : 1210
}
```

Example:

Using a non-existing collection

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/import?collection=products&type=documents <<
EOF
{ "name": "test" }
EOF

HTTP/1.1 404 Not Found
content-type: application/json; charset=utf-8

{
  "error" : true,
  "errorMessage" : "collection 'products' not found",
  "code" : 404,
  "errorNum" : 1203
}
```

Example:

Using a malformed body

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/import?collection=products&type=list <<EOF
{ }
EOF

HTTP/1.1 400 Bad Request
content-type: application/json; charset=utf-8

{
  "error" : true,
  "errorMessage" : "expecting a JSON array in the request",
  "code" : 400,
  "errorNum" : 400
}
```

Return Codes

- *201*: is returned if all documents could be imported successfully.
- *400*: is returned if `type` contains an invalid value, no `collection` is specified, the documents are incorrectly encoded, or the request is malformed.
- *404*: is returned if `collection` or the `_from` or `_to` attributes of an imported edge refer to an unknown collection.
- *409*: is returned if the import would trigger a unique key violation and `complete` is set to `true`.
- *500*: is returned if the server cannot auto-generate a document key (out of keys error) for a document with no user-defined key.
- *501*: The server will respond with *HTTP 501* if this API is called on a cluster coordinator.

Examples

Importing documents with heterogenous attributes from a JSON array

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/import?
collection=products&type=list <<EOF
[
  {
    "_key" : "abc",
    "value1" : 25,
    "value2" : "test",
    "allowed" : true
  },
  {
    "_key" : "foo",
    "name" : "baz"
  },
  {
    "name" : {
      "detailed" : "detailed name",
      "short" : "short name"
    }
  }
]
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8
```

show response body

Importing documents from individual JSON lines

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/import?
collection=products&type=documents <<EOF
{ "_key": "abc", "value1": 25, "value2": "test","allowed": true }
{ "_key": "foo", "name": "baz" }

{ "name": { "detailed": "detailed name", "short": "short name" } }

EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8
```

show response body

Using the auto type detection

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/import?
collection=products&type=auto <<EOF
[
  {
    "_key" : "abc",
    "value1" : 25,
    "value2" : "test",
    "allowed" : true
  },
  {
    "_key" : "foo",
    "name" : "baz"
  },
  {
    "name" : {
      "detailed" : "detailed name",
      "short" : "short name"
    }
  }
]
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8
```

show response body

Importing into an edge collection, with attributes `_from`, `_to` and `name`

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/import?
collection=links&type=documents <<EOF
{ "_from": "products/123", "_to": "products/234" }
{"_from": "products/332", "_to": "products/abc",   "name": "other name" }
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8
```

show response body

Importing into an edge collection, omitting `_from` or `_to`

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/import?
collection=links&type=list&details=true <<EOF
[
  {
    "name" : "some name"
  }
]
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8
```

show response body

Violating a unique constraint, but allow partial imports

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/import?
collection=products&type=documents&details=true <<EOF
{ "_key": "abc", "value1": 25, "value2": "test" }
{ "_key": "abc", "value1": "bar", "value2": "baz" }
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8
```

show response body

Violating a unique constraint, not allowing partial imports

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/import?
collection=products&type=documents&complete=true <<EOF
{ "_key": "abc", "value1": 25, "value2": "test" }
{ "_key": "abc", "value1": "bar", "value2": "baz" }
EOF

HTTP/1.1 409 Conflict
content-type: application/json; charset=utf-8
```

show response body

Using a non-existing collection

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/import?
collection=products&type=documents <<EOF
{ "name": "test" }
EOF

HTTP/1.1 404 Not Found
content-type: application/json; charset=utf-8
```

show response body

Using a malformed body

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/import?
collection=products&type=list <<EOF
{ }
EOF

HTTP/1.1 400 Bad Request
content-type: application/json; charset=utf-8
```

show response body

Importing Self-Contained JSON Documents

This import method allows uploading self-contained JSON documents. The documents must be uploaded in the body of the HTTP POST request. Each line of the body will be interpreted as one stand-alone document. Empty lines in the body are allowed but will be skipped. Using this format, the documents are imported line-wise.

Example input data: { "_key": "key1", ... } { "_key": "key2", ... } ...

To use this method, the *type* query parameter should be set to *documents*.

It is also possible to upload self-contained JSON documents that are embedded into a JSON array. Each element from the array will be treated as a document and be imported.

Example input data for this case:

```
[
  { "_key": "key1", ... },
  { "_key": "key2", ... },
  ...
]
```

This format does not require each document to be on a separate line, and any whitespace in the JSON data is allowed. It can be used to import a JSON-formatted result array (e.g. from arangosh) back into ArangoDB. Using this format requires ArangoDB to parse the complete array and keep it in memory for the duration of the import. This might be more resource-intensive than the line-wise processing.

To use this method, the *type* query parameter should be set to *array*.

Setting the *type* query parameter to *auto* will make the server auto-detect whether the data are line-wise JSON documents (type = documents) or a JSON array (type = array).

Examples

```
curl --data-binary @- -X POST --dump - "http://localhost:8529/_api/import?type=documents&collection=test"
{ "name" : "test", "gender" : "male", "age" : 39 }
{ "type" : "bird", "name" : "robin" }

HTTP/1.1 201 Created
Server: ArangoDB
Connection: Keep-Alive
Content-type: application/json; charset=utf-8

{"error":false,"created":2,"empty":0,"errors":0}
```

The server will respond with an HTTP 201 if everything went well. The number of documents imported will be returned in the *created* attribute of the response. If any documents were skipped or incorrectly formatted, this will be returned in the *errors* attribute. There will also be an attribute *empty* in the response, which will contain a value of *0*.

If the *details* parameter was set to *true* in the request, the response will also contain an attribute *details* which is an array of details about errors that occurred on the server side during the import. This array might be empty if no errors occurred.

Importing Headers and Values

When using this type of import, the attribute names of the documents to be imported are specified separate from the actual document value data. The first line of the HTTP POST request body must be a JSON array containing the attribute names for the documents that follow. The following lines are interpreted as the document data. Each document must be a JSON array of values. No attribute names are needed or allowed in this data section.

Examples

```
curl --data-binary @- -X POST --dump - "http://localhost:8529/_api/import?collection=test"
[ "firstName", "lastName", "age", "gender" ]
[ "Joe", "Public", 42, "male" ]
[ "Jane", "Doe", 31, "female" ]

HTTP/1.1 201 Created
Server: ArangoDB
Connection: Keep-Alive
Content-type: application/json; charset=utf-8

{"error":false,"created":2,"empty":0,"errors":0}
```

The server will again respond with an HTTP 201 if everything went well. The number of documents imported will be returned in the *created* attribute of the response. If any documents were skipped or incorrectly formatted, this will be returned in the *errors* attribute. The number of empty lines in the input file will be returned in the *empty* attribute.

If the *details* parameter was set to *true* in the request, the response will also contain an attribute *details* which is an array of details about errors that occurred on the server side during the import. This array might be empty if no errors occurred.

Importing into Edge Collections

Please note that when importing documents into an edge collection, it is mandatory that all imported documents contain the *_from* and *_to* attributes, and that these contain references to existing collections.

HTTP Interface for Batch Requests

Clients normally send individual operations to ArangoDB in individual HTTP requests. This is straightforward and simple, but has the disadvantage that the network overhead can be significant if many small requests are issued in a row.

To mitigate this problem, ArangoDB offers a batch request API that clients can use to send multiple operations in one batch to ArangoDB. This method is especially useful when the client has to send many HTTP requests with a small body/payload and the individual request results do not depend on each other.

Clients can use ArangoDB's batch API by issuing a multipart HTTP POST request to the URL */_api/batch* handler. The handler will accept the request if the Content-type is *multipart/form-data* and a boundary string is specified. ArangoDB will then decompose the batch request into its individual parts using this boundary. This also means that the boundary string itself must not be contained in any of the parts. When ArangoDB has split the multipart request into its individual parts, it will process all parts sequentially as if it were a standalone request. When all parts are processed, ArangoDB will generate a multipart HTTP response that contains one part for each part operation result. For example, if you send a multipart request with 5 parts, ArangoDB will send back a multipart response with 5 parts as well.

The server expects each part message to start with exactly the following "header":

```
Content-type: application/x-arango-batchpart
```

You can optionally specify a *Content-Id* "header" to uniquely identify each part message. The server will return the *Content-Id* in its response if it is specified. Otherwise, the server will not send a Content-Id "header" back. The server will not validate the uniqueness of the Content-Id. After the mandatory *Content-type* and the optional *Content-Id* header, two Windows line breaks (i.e. *\r\n\r\n*) must follow. Any deviation of this structure might lead to the part being rejected or incorrectly interpreted. The part request payload, formatted as a regular HTTP request, must follow the two Windows line breaks literal directly.

Note that the literal *Content-type: application/x-arango-batchpart* technically is the header of the MIME part, and the HTTP request (including its headers) is the body part of the MIME part.

An actual part request should start with the HTTP method, the called URL, and the HTTP protocol version as usual, followed by arbitrary HTTP headers. Its body should follow after the usual *\r\n\r\n* literal. Part requests are therefore regular HTTP requests, only embedded inside a multipart message.

The following example will send a batch with 3 individual document creation operations. The boundary used in this example is *XXXsubpartXXX*.

Examples

```
> curl -X POST --data-binary @- --header "Content-type: multipart/form-data; boundary=XXXsubpartXXX" http://localhost
:8529/_api/batch
--XXXsubpartXXX
Content-type: application/x-arango-batchpart
Content-Id: 1

POST /_api/document?collection=xyz HTTP/1.1

{"a":1,"b":2,"c":3}
--XXXsubpartXXX
Content-type: application/x-arango-batchpart
Content-Id: 2

POST /_api/document?collection=xyz HTTP/1.1

{"a":1,"b":2,"c":3,"d":4}
--XXXsubpartXXX
Content-type: application/x-arango-batchpart
Content-Id: 3

POST /_api/document?collection=xyz HTTP/1.1

{"a":1,"b":2,"c":3,"d":4,"e":5}
--XXXsubpartXXX--
```

The server will then respond with one multipart message, containing the overall status and the individual results for the part operations. The overall status should be 200 except there was an error while inspecting and processing the multipart message. The overall status therefore does not indicate the success of each part operation, but only indicates whether the multipart message could be handled successfully.

Each part operation will return its own status value. As the part operation results are regular HTTP responses (just included in one multipart response), the part operation status is returned as a HTTP status code. The status codes of the part operations are exactly the same as if you called the individual operations standalone. Each part operation might also return arbitrary HTTP headers and a body/payload:

Examples

```
HTTP/1.1 200 OK
Connection: Keep-Alive
Content-type: multipart/form-data; boundary=XXXsubpartXXX
Content-length: 1055

--XXXsubpartXXX
Content-type: application/x-arango-batchpart
Content-Id: 1

HTTP/1.1 202 Accepted
Content-type: application/json; charset=utf-8
Etag: "9514299"
Content-length: 53

{"error":false,"_id":"xyz/9514299","_key":"9514299","_rev":"9514299"}
--XXXsubpartXXX
Content-type: application/x-arango-batchpart
Content-Id: 2

HTTP/1.1 202 Accepted
Content-type: application/json; charset=utf-8
Etag: "9579835"
Content-length: 53

{"error":false,"_id":"xyz/9579835","_key":"9579835","_rev":"9579835"}
--XXXsubpartXXX
Content-type: application/x-arango-batchpart
Content-Id: 3

HTTP/1.1 202 Accepted
Content-type: application/json; charset=utf-8
Etag: "9645371"
Content-length: 53

{"error":false,"_id":"xyz/9645371","_key":"9645371","_rev":"9645371"}
--XXXsubpartXXX--
```

In the above example, the server returned an overall status code of 200, and each part response contains its own status value (202 in the example):

When constructing the multipart HTTP response, the server will use the same boundary that the client supplied. If any of the part responses has a status code of 400 or greater, the server will also return an HTTP header *x-arango-errors* containing the overall number of part requests that produced errors:

Examples

```
> curl -X POST --data-binary @- --header "Content-type: multipart/form-data; boundary=XXXsubpartXXX" http://localhost
:8529/_api/batch
--XXXsubpartXXX
Content-type: application/x-arango-batchpart

POST /_api/document?collection=nonexisting

{"a":1,"b":2,"c":3}
--XXXsubpartXXX
Content-type: application/x-arango-batchpart

POST /_api/document?collection=xyz

{"a":1,"b":2,"c":3,"d":4}
--XXXsubpartXXX--
```

In this example, the overall response code is 200, but as some of the part request failed (with status code 404), the *x-arango-errors* header of the overall response is *1*:

Examples

```
HTTP/1.1 200 OK
x-arango-errors: 1
Content-type: multipart/form-data; boundary=XXXsubpartXXX
Content-length: 711

--XXXsubpartXXX
Content-type: application/x-arango-batchpart

HTTP/1.1 404 Not Found
Content-type: application/json; charset=utf-8
Content-length: 111

{"error":true,"code":404,"errorNum":1203,"errorMessage":"collection \/_api\/collection\/nonexisting not found"}
--XXXsubpartXXX
Content-type: application/x-arango-batchpart

HTTP/1.1 202 Accepted
Content-type: application/json; charset=utf-8
Etag: "9841979"
Content-length: 53

{"error":false,"_id":"xyz/9841979","_key":"9841979","_rev":"9841979"}
--XXXsubpartXXX--
```

Please note that the database used for all part operations of a batch request is determined by scanning the original URL (the URL that contains */_api/batch*). It is not possible to override the database name in part operations of a batch. When doing so, any other database name used in a batch part will be ignored.

executes a batch request

executes a batch request

`POST /_api/batch`

Post Body

body (required): The multipart batch request, consisting of the envelope and the individual batch parts.

free style json body

Executes a batch request. A batch request can contain any number of other requests that can be sent to ArangoDB in isolation. The benefit of using batch requests is that batching requests requires less client/server roundtrips than when sending isolated requests.

All parts of a batch request are executed serially on the server. The server will return the results of all parts in a single response when all parts are finished.

Technically, a batch request is a multipart HTTP request, with content-type `multipart/form-data`. A batch request consists of an envelope and the individual batch part actions. Batch part actions are "regular" HTTP requests, including full header and an optional body. Multiple batch parts are separated by a boundary identifier. The boundary identifier is declared in the batch envelope. The MIME content-type for each individual batch part must be `application/x-arango-batchpart`.

Please note that when constructing the individual batch parts, you must use CRLF (`\ \`) as the line terminator as in regular HTTP messages.

The response sent by the server will be an `HTTP 200` response, with an optional error summary header `x-arango-errors`. This header contains the number of batch part operations that failed with an HTTP error code of at least 400. This header is only present in the response if the number of errors is greater than zero.

The response sent by the server is a multipart response, too. It contains the individual HTTP responses for all batch parts, including the full HTTP result header (with status code and other potential headers) and an optional result body. The individual batch parts in the result are seperated using the same boundary value as specified in the request.

The order of batch parts in the response will be the same as in the original client request. Client can additionally use the `Content-Id` MIME header in a batch part to define an individual id for each batch part. The server will return this id is the batch part responses, too.

Example:

Sending a batch request with five batch parts:

- GET /_api/version
- DELETE /_api/collection/products
- POST /_api/collection/products
- GET /_api/collection/products/figures
- DELETE /_api/collection/products

The boundary (`SomeBoundaryValue`) is passed to the server in the HTTP `Content-Type` HTTP header. *Please note the reply is not displayed all accurate.*

```
shell> curl -X POST --header 'Content-Type: multipart/form-data; boundary=SomeBoundaryValue' --data-binary @- --dump
- http://localhost:8529/_api/batch <<EOF
--SomeBoundaryValue
Content-Type: application/x-arango-batchpart
Content-Id: myId1

GET /_api/version HTTP/1.1

--SomeBoundaryValue
Content-Type: application/x-arango-batchpart
Content-Id: myId2

DELETE /_api/collection/products HTTP/1.1

--SomeBoundaryValue
Content-Type: application/x-arango-batchpart
Content-Id: someId

POST /_api/collection/products HTTP/1.1

{"name": "products" }

--SomeBoundaryValue
Content-Type: application/x-arango-batchpart
Content-Id: nextId

GET /_api/collection/products/figures HTTP/1.1

--SomeBoundaryValue
Content-Type: application/x-arango-batchpart
Content-Id: otherId

DELETE /_api/collection/products HTTP/1.1
```

```
--SomeBoundaryValue--

EOF

HTTP/1.1 200 OK
content-type: multipart/form-data; boundary=SomeBoundaryValue
x-arango-errors: 1

"--SomeBoundaryValue
Content-Type: application/x-arango-batchpart
Content-Id: myId1

HTTP/1.1 200 OK
Server:
Connection:
Content-Type: application/json; charset=utf-8
Content-Length: 60

{\"server\":\"arango\",\"version\":\"3.1.18\",\"license\":\"community\"}
--SomeBoundaryValue
Content-Type: application/x-arango-batchpart
Content-Id: myId2

HTTP/1.1 404 Not Found
Server:
Content-Type: application/json; charset=utf-8
Connection:
Content-Length: 88

{\"error\":true,\"code\":404,\"errorNum\":1203,\"errorMessage\":\"unknown collection 'products'\"}
--SomeBoundaryValue
Content-Type: application/x-arango-batchpart
Content-Id: someId

HTTP/1.1 200 OK
Server:
Location: /_db/_system/_api/collection/products
Content-Type: application/json; charset=utf-8
Connection:
Content-Length: 132

{\"id\":\"9101\",\"name\":\"products\",\"waitForSync\":false,\"isVolatile\":false,\"isSystem\":false,\"status\":3,\"t
ype\":2,\"error\":false,\"code\":200}
--SomeBoundaryValue
Content-Type: application/x-arango-batchpart
Content-Id: nextId

HTTP/1.1 200 OK
Server:
Location: /_db/_system/_api/collection/products/figures
Content-Type: application/json; charset=utf-8
Connection:
Content-Length: 747

{\"id\":\"9101\",\"name\":\"products\",\"isSystem\":false,\"doCompact\":true,\"isVolatile\":false,\"journalSize\":335
54432,\"keyOptions\":{\"type\":\"traditional\",\"allowUserKeys\":true},\"waitForSync\":false,\"indexBuckets\":8,\"cou
nt\":0,\"figures\":{\"indexes\":{\"count\":1,\"size\":32128},\"lastTick\":0,\"uncollectedLogfileEntries\":0,\"compact
ionStatus\":{\"message\":\"compaction not yet started\",\"time\":\"2017-04-18T15:52:31Z\"},\"readCache\":{\"count\":0
,\"size\":48192,\"total\":8912896},\"documentReferences\":0,\"waitingFor\":\"-\",\"alive\":{\"count\":0,\"size\":0},\
"dead\":{\"count\":0,\"size\":0,\"deletion\":0},\"datafiles\":{\"count\":0,\"fileSize\":0},\"journals\":{\"count\":0,
\"fileSize\":0},\"compactors\":{\"count\":0,\"fileSize\":0},\"revisions\":{\"count\":0,\"size\":48192}},\"status\":3,
\"type\":2,\"error\":false,\"code\":200}
--SomeBoundaryValue
Content-Type: application/x-arango-batchpart
Content-Id: otherId

HTTP/1.1 200 OK
Server:
Content-Type: application/json; charset=utf-8
Connection:
Content-Length: 38

{\"id\":\"9101\",\"error\":false,\"code\":200}
--SomeBoundaryValue--"
```

Example:

Sending a batch request, setting the boundary implicitly (the server will in this case try to find the boundary at the beginning of the request body).

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/batch <<EOF
--SomeBoundaryValue
Content-Type: application/x-arango-batchpart

DELETE /_api/collection/notexisting1 HTTP/1.1

--SomeBoundaryValue
Content-Type: application/x-arango-batchpart

DELETE _api/collection/notexisting2 HTTP/1.1
--SomeBoundaryValue--

EOF

HTTP/1.1 200 OK
content-type:
x-arango-errors: 2

"--SomeBoundaryValue
Content-Type: application/x-arango-batchpart

HTTP/1.1 404 Not Found
Server:
Content-Type: application/json; charset=utf-8
Connection:
Content-Length: 92

{\"error\":true,\"code\":404,\"errorNum\":1203,\"errorMessage\":\"unknown collection 'notexisting1'\"}
--SomeBoundaryValue
Content-Type: application/x-arango-batchpart

HTTP/1.1 404 Not Found
Server:
Content-Type: application/json; charset=utf-8
Connection:
Content-Length: 101

{\"error\":true,\"code\":404,\"errorNum\":404,\"errorMessage\":\"unknown path '_api/collection/notexisting2'\"}
--SomeBoundaryValue--"
```

Return Codes

- *200*: is returned if the batch was received successfully. HTTP 200 is returned even if one or multiple batch part actions failed.
- *400*: is returned if the batch envelope is malformed or incorrectly formatted. This code will also be returned if the content-type of the overall batch request or the individual MIME parts is not as expected.
- *405*: is returned when an invalid HTTP method is used.

Examples

Sending a batch request with five batch parts:

- GET /_api/version
- DELETE /_api/collection/products
- POST /_api/collection/products
- GET /_api/collection/products/figures
- DELETE /_api/collection/products The boundary (`SomeBoundaryValue`) is passed to the server in the HTTP `Content-Type` HTTP header. *Please note the reply is not displayed all accurate.*

```
shell> curl -X POST --header 'Content-Type: multipart/form-data;
boundary=SomeBoundaryValue' --data-binary @- --dump - http://localhost:8529/_api/batch
<<EOF
--SomeBoundaryValue
```

```
Content-Type: application/x-arango-batchpart
Content-Id: myId1

GET /_api/version HTTP/1.1
--SomeBoundaryValue
Content-Type: application/x-arango-batchpart
Content-Id: myId2

DELETE /_api/collection/products HTTP/1.1
--SomeBoundaryValue
Content-Type: application/x-arango-batchpart
Content-Id: someId

POST /_api/collection/products HTTP/1.1

{"name": "products" }
--SomeBoundaryValue
Content-Type: application/x-arango-batchpart
Content-Id: nextId

GET /_api/collection/products/figures HTTP/1.1
--SomeBoundaryValue
Content-Type: application/x-arango-batchpart
Content-Id: otherId

DELETE /_api/collection/products HTTP/1.1
--SomeBoundaryValue--

EOF

HTTP/1.1 200 OK
content-type: multipart/form-data; boundary=SomeBoundaryValue
x-arango-errors: 1

"--SomeBoundaryValue\r\nContent-Type: application/x-arango-batchpart\r\nContent-Id:
myId1\r\n\r\nHTTP/1.1 200 OK\r\nServer: \r\nConnection: \r\nContent-Type:
application/json; charset=utf-8\r\nContent-Length:
60\r\n\r\n{\"server\":\"arango\",\"version\":\"3.1.18\",\"license\":\"community\"}\r\n--
SomeBoundaryValue\r\nContent-Type: application/x-arango-batchpart\r\nContent-Id:
myId2\r\n\r\nHTTP/1.1 404 Not Found\r\nServer: \r\nContent-Type: application/json;
charset=utf-8\r\nConnection: \r\nContent-Length:
88\r\n\r\n{\"error\":true,\"code\":404,\"errorNum\":1203,\"errorMessage\":\"unknown
collection 'products'\"}\r\n--SomeBoundaryValue\r\nContent-Type: application/x-arango-
batchpart\r\nContent-Id: someId\r\n\r\nHTTP/1.1 200 OK\r\nServer: \r\nLocation:
/_db/_system/_api/collection/products\r\nContent-Type: application/json; charset=utf-
8\r\nConnection: \r\nContent-Length:
132\r\n\r\n{\"id\":\"9101\",\"name\":\"products\",\"waitForSync\":false,\"isVolatile\":f
alse,\"isSystem\":false,\"status\":3,\"type\":2,\"error\":false,\"code\":200}\r\n--
SomeBoundaryValue\r\nContent-Type: application/x-arango-batchpart\r\nContent-Id:
nextId\r\n\r\nHTTP/1.1 200 OK\r\nServer: \r\nLocation:
/_db/_system/_api/collection/products/figures\r\nContent-Type: application/json;
charset=utf-8\r\nConnection: \r\nContent-Length:
747\r\n\r\n{\"id\":\"9101\",\"name\":\"products\",\"isSystem\":false,\"doCompact\":true,
\"isVolatile\":false,\"journalSize\":33554432,\"keyOptions\":
{\"type\":\"traditional\",\"allowUserKeys\":true},\"waitForSync\":false,\"indexBuckets\"
:8,\"count\":0,\"figures\":{\"indexes\":
```

```
{\"count\":1,\"size\":32128},\"lastTick\":0,\"uncollectedLogfileEntries\":0,\"compaction
Status\":{\"message\":\"compaction not yet started\",\"time\":\"2017-04-
18T15:52:31Z\"},\"readCache\":
{\"count\":0,\"size\":48192,\"total\":8912896},\"documentReferences\":0,\"waitingFor\":\
"-\",\"alive\":{\"count\":0,\"size\":0},\"dead\":
{\"count\":0,\"size\":0,\"deletion\":0},\"datafiles\":
{\"count\":0,\"fileSize\":0},\"journals\":{\"count\":0,\"fileSize\":0},\"compactors\":
{\"count\":0,\"fileSize\":0},\"revisions\":
{\"count\":0,\"size\":48192}},\"status\":3,\"type\":2,\"error\":false,\"code\":200}\r\n-
-SomeBoundaryValue\r\nContent-Type: application/x-arango-batchpart\r\nContent-Id:
otherId\r\n\r\nHTTP/1.1 200 OK\r\nServer: \r\nContent-Type: application/json;
charset=utf-8\r\nConnection: \r\nContent-Length:
38\r\n\r\n{\"id\":\"9101\",\"error\":false,\"code\":200}\r\n--SomeBoundaryValue--"
```

Sending a batch request, setting the boundary implicitly (the server will in this case try to find the boundary at the beginning of the request body).

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/batch <<EOF
--SomeBoundaryValue
Content-Type: application/x-arango-batchpart

DELETE /_api/collection/notexisting1 HTTP/1.1
--SomeBoundaryValue
Content-Type: application/x-arango-batchpart

DELETE _api/collection/notexisting2 HTTP/1.1
--SomeBoundaryValue--

EOF

HTTP/1.1 200 OK
content-type:
x-arango-errors: 2

"--SomeBoundaryValue\r\nContent-Type: application/x-arango-batchpart\r\n\r\nHTTP/1.1 404
Not Found\r\nServer: \r\nContent-Type: application/json; charset=utf-8\r\nConnection:
\r\nContent-Length:
92\r\n\r\n{\"error\":true,\"code\":404,\"errorNum\":1203,\"errorMessage\":\"unknown
collection 'notexisting1'\"}\r\n--SomeBoundaryValue\r\nContent-Type: application/x-
arango-batchpart\r\n\r\nHTTP/1.1 404 Not Found\r\nServer: \r\nContent-Type:
application/json; charset=utf-8\r\nConnection: \r\nContent-Length:
101\r\n\r\n{\"error\":true,\"code\":404,\"errorNum\":404,\"errorMessage\":\"unknown path
'_api/collection/notexisting2'\"}\r\n--SomeBoundaryValue--"
```

Deletes an existing dump batch

handle a dump batch command

`DELETE /_api/replication/batch/{id}` **Note**: These calls are uninteresting to users.

Path Parameters

- *id* (required): The id of the batch.

Deletes the existing dump batch, allowing compaction and cleanup to resume.

Note: on a coordinator, this request must have the query parameter *DBserver* which must be an ID of a DBserver. The very same request is forwarded synchronously to that DBserver. It is an error if this attribute is not bound in the coordinator case.

Return Codes

- *204*: is returned if the batch was deleted successfully.
- *400*: is returned if the batch was not found.
- *405*: is returned when an invalid HTTP method is used.

Create new dump batch

handle a dump batch command

`POST /_api/replication/batch` **Note**: These calls are uninteresting to users.

A json post document with these Properties is required:

- **ttl**: the time-to-live for the new batch (in seconds) A JSON object with the batch configuration.

Creates a new dump batch and returns the batch's id.

The response is a JSON object with the following attributes:

- *id*: the id of the batch

Note: on a coordinator, this request must have the query parameter *DBserver* which must be an ID of a DBserver. The very same request is forwarded synchronously to that DBserver. It is an error if this attribute is not bound in the coordinator case.

Return Codes

- *204*: is returned if the batch was created successfully.
- *400*: is returned if the ttl value is invalid or if *DBserver* attribute is not specified or illegal on a coordinator.
- *405*: is returned when an invalid HTTP method is used.

Prolong existing dump batch

handle a dump batch command

`PUT /_api/replication/batch/{id}` **Note**: These calls are uninteresting to users.

A json post document with these Properties is required:

- **ttl**: the time-to-live for the new batch (in seconds)

Extends the ttl of an existing dump batch, using the batch's id and the provided ttl value.

If the batch's ttl can be extended successfully, the response is empty.

Note: on a coordinator, this request must have the query parameter *DBserver* which must be an ID of a DBserver. The very same request is forwarded synchronously to that DBserver. It is an error if this attribute is not bound in the coordinator case.

Path Parameters

- *id* (required): The id of the batch.

Return Codes

- *204*: is returned if the batch's ttl was extended successfully.
- *400*: is returned if the ttl value is invalid or the batch was not found.
- *405*: is returned when an invalid HTTP method is used.

HTTP Interface for Exporting Documents

Create export cursor

export all documents from a collection, using a cursor

`POST /_api/export`

A json post document with these Properties is required:

- **count**: boolean flag that indicates whether the number of documents in the result set should be returned in the "count" attribute of the result (optional). Calculating the "count" attribute might in the future have a performance impact so this option is turned off by default, and "count" is only returned when requested.
- **restrict**:
  - **fields** (string): Contains an array of attribute names to *include* or *exclude*. Matching of attribute names for *inclusion* or *exclusion* will be done on the top level only. Specifying names of nested attributes is not supported at the moment.
  - **type**: has to be be set to either *include* or *exclude* depending on which you want to use
- **batchSize**: maximum number of result documents to be transferred from the server to the client in one roundtrip (optional). If this attribute is not set, a server-controlled default value will be used.
- **flush**: if set to *true*, a WAL flush operation will be executed prior to the export. The flush operation will start copying documents from the WAL to the collection's datafiles. There will be an additional wait time of up to *flushWait* seconds after the flush to allow the WAL collector to change the adjusted document meta-data to point into the datafiles, too. The default value is *false* (i.e. no flush) so most recently inserted or updated documents from the collection might be missing in the export.
- **flushWait**: maximum wait time in seconds after a flush operation. The default value is 10. This option only has an effect when *flush* is set to *true*.
- **limit**: an optional limit value, determining the maximum number of documents to be included in the cursor. Omitting the *limit* attribute or setting it to 0 will lead to no limit being used. If a limit is used, it is undefined which documents from the collection will be included in the export and which will be excluded. This is because there is no natural order of documents in a collection.
- **ttl**: an optional time-to-live for the cursor (in seconds). The cursor will be removed on the server automatically after the specified amount of time. This is useful to ensure garbage collection of cursors that are not fully fetched by clients. If not set, a server-defined value will be used.

A call to this method creates a cursor containing all documents in the specified collection. In contrast to other data-producing APIs, the internal data structures produced by the export API are more lightweight, so it is the preferred way to retrieve all documents from a collection.

Documents are returned in a similar manner as in the `/_api/cursor` REST API. If all documents of the collection fit into the first batch, then no cursor will be created, and the result object's *hasMore* attribute will be set to *false*. If not all documents fit into the first batch, then the result object's *hasMore* attribute will be set to *true*, and the *id* attribute of the result will contain a cursor id.

The order in which the documents are returned is not specified.

By default, only those documents from the collection will be returned that are stored in the collection's datafiles. Documents that are present in the write-ahead log (WAL) at the time the export is run will not be exported.

To export these documents as well, the caller can issue a WAL flush request before calling the export API or set the *flush* attribute. Setting the *flush* option will trigger a WAL flush before the export so documents get copied from the WAL to the collection datafiles.

If the result set can be created by the server, the server will respond with *HTTP 201*. The body of the response will contain a JSON object with the result set.

The returned JSON object has the following properties:

- *error*: boolean flag to indicate that an error occurred (*false* in this case)
- *code*: the HTTP status code
- *result*: an array of result documents (might be empty if the collection was empty)
- *hasMore*: a boolean indicator whether there are more results available for the cursor on the server
- *count*: the total number of result documents available (only available if the query was executed with the *count* attribute set)

- *id*: id of temporary cursor created on the server (optional, see above)

If the JSON representation is malformed or the query specification is missing from the request, the server will respond with *HTTP 400*.

The body of the response will contain a JSON object with additional error details. The object has the following attributes:

- *error*: boolean flag to indicate that an error occurred (*true* in this case)
- *code*: the HTTP status code
- *errorNum*: the server error number
- *errorMessage*: a descriptive error message

Clients should always delete an export cursor result as early as possible because a lingering export cursor will prevent the underlying collection from being compacted or unloaded. By default, unused cursors will be deleted automatically after a server-defined idle time, and clients can adjust this idle time by setting the *ttl* value.

Note: this API is currently not supported on cluster coordinators.

HTTP Interface for Indexes

Indexes

This is an introduction to ArangoDB's HTTP interface for indexes in general. There are special sections for various index types.

Index

Indexes are used to allow fast access to documents. For each collection there is always the primary index which is a hash index for the document key (_key attribute). This index cannot be dropped or changed. edge collections will also have an automatically created edges index, which cannot be modified. This index provides quick access to documents via the `_from` and `_to` attributes.

Most user-land indexes can be created by defining the names of the attributes which should be indexed. Some index types allow indexing just one attribute (e.g. fulltext index) whereas other index types allow indexing multiple attributes.

Using the system attribute `_id` in user-defined indexes is not supported by any index type.

Index Handle

An index handle uniquely identifies an index in the database. It is a string and consists of a collection name and an index identifier separated by /. Geo Index: A geo index is used to find places on the surface of the earth fast. Hash Index: A hash index is used to find documents based on examples. A hash index can be created for one or multiple document attributes. A hash index will only be used by queries if all indexed attributes are present in the example or search query, and if all attributes are compared using the equality (== operator). That means the hash index does not support range queries.

If the index is declared unique, then access to the indexed attributes should be fast. The performance degrades if the indexed attribute(s) contain(s) only very few distinct values.

Edges Index

An edges index is automatically created for edge collections. It contains connections between vertex documents and is invoked when the connecting edges of a vertex are queried. There is no way to explicitly create or delete edge indexes.

Skiplist Index

A skiplist is a sorted index that can be used to find individual documents or ranges of documents.

Persistent Index

A persistent index is a sorted index that can be used for finding individual documents or ranges of documents. In constrast to the other indexes, the contents of a persistent index are stored on disk and thus do not need to be rebuilt in memory from the documents when the collection is loaded.

Fulltext Index:

A fulltext index can be used to find words, or prefixes of words inside documents. A fulltext index can be set on one attribute only, and will index all words contained in documents that have a textual value in this attribute. Only words with a (specifiable) minimum length are indexed. Word tokenization is done using the word boundary analysis provided by libicu, which is taking into account the selected language provided at server start. Words are indexed in their lower-cased form. The index supports complete match queries (full words) and prefix queries.

The basic operations (create, read, update, delete) for documents are mapped to the standard HTTP methods (*POST*, *GET*, *PUT*, *DELETE*).

Address of an Index

All indexes in ArangoDB have an unique handle. This index handle identifies an index and is managed by ArangoDB. All indexes are found under the URI

```
http://server:port/_api/index/index-handle
```

For example: Assume that the index handle is *demo/63563528* then the URL of that index is:

```
http://localhost:8529/_api/index/demo/63563528
```

Working with Indexes using HTTP

Read index

returns an index

`GET /_api/index/{index-handle}`

Path Parameters

- *index-handle* (required): The index-handle.

The result is an object describing the index. It has at least the following attributes:

- *id*: the identifier of the index
- *type*: the index type

All other attributes are type-dependent. For example, some indexes provide *unique* or *sparse* flags, whereas others don't. Some indexes also provide a selectivity estimate in the *selectivityEstimate* attribute of the result.

Example:

```
shell> curl --dump - http://localhost:8529/_api/index/products/0

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "id" : "products/0",
  "type" : "primary",
  "fields" : [
    "_key"
  ],
  "selectivityEstimate" : 1,
  "unique" : true,
  "sparse" : false,
  "error" : false,
  "code" : 200
}
```

Return Codes

- *200*: If the index exists, then a *HTTP 200* is returned.
- *404*: If the index does not exist, then a *HTTP 404* is returned.

Examples

```
shell> curl --dump - http://localhost:8529/_api/index/products/0

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Create index

creates an index

`POST /_api/index#general`

Query Parameters

- *collection* (required): The collection name.

Post Body

index-details (required):

free style json body

NOTE Swagger examples won't work due to the anchor.

Creates a new index in the collection *collection*. Expects an object containing the index details.

The type of the index to be created must specified in the *type* attribute of the index details. Depending on the index type, additional other attributes may need to specified in the request in order to create the index.

Indexes require the to be indexed attribute(s) in the *fields* attribute of the index details. Depending on the index type, a single attribute or multiple attributes can be indexed. In the latter case, an array of strings is expected.

Indexing the system attribute *_id* is not supported for user-defined indexes. Manually creating an index using *_id* as an index attribute will fail with an error.

Some indexes can be created as unique or non-unique variants. Uniqueness can be controlled for most indexes by specifying the *unique* flag in the index details. Setting it to *true* will create a unique index. Setting it to *false* or omitting the *unique* attribute will create a non-unique index.

Note: The following index types do not support uniqueness, and using the *unique* attribute with these types may lead to an error:

- geo indexes
- fulltext indexes

Note: Unique indexes on non-shard keys are not supported in a cluster.

Hash, skiplist and persistent indexes can optionally be created in a sparse variant. A sparse index will be created if the *sparse* attribute in the index details is set to *true*. Sparse indexes do not index documents for which any of the index attributes is either not set or is *null*.

Return Codes

- *200*: If the index already exists, then an *HTTP 200* is returned.
- *201*: If the index does not already exist and could be created, then an *HTTP 201* is returned.
- *400*: If an invalid index description is posted or attributes are used that the target index will not support, then an *HTTP 400* is returned.
- *404*: If *collection* is unknown, then an *HTTP 404* is returned.

Delete index

deletes an index

`DELETE /_api/index/{index-handle}`

Path Parameters

- *index-handle* (required): The index handle.

Deletes an index with *index-handle*.

Example:

```
shell> curl -X DELETE --dump - http://localhost:8529/_api/index/products/10475

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "id" : "products/10475",
  "error" : false,
  "code" : 200
}
```

Return Codes

- *200*: If the index could be deleted, then an *HTTP 200* is returned.
- *404*: If the *index-handle* is unknown, then an *HTTP 404* is returned.

Examples

```
shell> curl -X DELETE --dump - http://localhost:8529/_api/index/products/10475

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Read all indexes of a collection

returns all indexes of a collection

`GET /_api/index`

Query Parameters

- *collection* (required): The collection name.

Returns an object with an attribute *indexes* containing an array of all index descriptions for the given collection. The same information is also available in the *identifiers* as an object with the index handles as keys.

Example:

Return information about all indexes

```
shell> curl --dump - http://localhost:8529/_api/index?collection=products

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "indexes" : [
    {
      "id" : "products/0",
      "type" : "primary",
      "fields" : [
        "_key"
      ],
      "selectivityEstimate" : 1,
      "unique" : true,
      "sparse" : false
    },
    {
      "id" : "products/10396",
      "type" : "hash",
      "fields" : [
        "name"
      ],
      "selectivityEstimate" : 1,
      "unique" : false,
      "sparse" : false
    },
    {
      "id" : "products/10399",
      "type" : "skiplist",
      "fields" : [
        "price"
      ],
      "unique" : false,
      "sparse" : true
    }
  ],
  "identifiers" : {
    "products/0" : {
      "id" : "products/0",
      "type" : "primary",
      "fields" : [
        "_key"
      ],
      "selectivityEstimate" : 1,
      "unique" : true,
      "sparse" : false
    },
    "products/10396" : {
      "id" : "products/10396",
      "type" : "hash",
      "fields" : [
        "name"
      ],
      "selectivityEstimate" : 1,
      "unique" : false,
      "sparse" : false
    },
    "products/10399" : {
      "id" : "products/10399",
      "type" : "skiplist",
      "fields" : [
        "price"
      ],
      "unique" : false,
      "sparse" : true
    }
  },
  "error" : false,
  "code" : 200
}
```

Return Codes

- *200*: returns a json object containing a list of indexes on that collection.

Examples

Return information about all indexes

```
shell> curl --dump - http://localhost:8529/_api/index?collection=products

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Working with Hash Indexes

If a suitable hash index exists, then */_api/simple/by-example* will use this index to execute a query-by-example.

Create hash index

creates a hash index

`POST /_api/index#hash`

Query Parameters

- *collection-name* (required): The collection name.

A json post document with these Properties is required:

- **fields** (string): an array of attribute paths.
- **unique**: if *true*, then create a unique index.
- **type**: must be equal to *"hash"*.
- **sparse**: if *true*, then create a sparse index.

NOTE Swagger examples won't work due to the anchor.

Creates a hash index for the collection *collection-name* if it does not already exist. The call expects an object containing the index details.

In a sparse index all documents will be excluded from the index that do not contain at least one of the specified index attributes (i.e. *fields*) or that have a value of *null* in any of the specified index attributes. Such documents will not be indexed, and not be taken into account for uniqueness checks if the *unique* flag is set.

In a non-sparse index, these documents will be indexed (for non-present indexed attributes, a value of *null* will be used) and will be taken into account for uniqueness checks if the *unique* flag is set.

Note: unique indexes on non-shard keys are not supported in a cluster.

Example:

Creating an unique constraint

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/index?collection=products <<EOF
{
  "type" : "hash",
  "unique" : true,
  "fields" : [
    "a",
    "b"
  ]
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8

{
  "id" : "products/10451",
  "type" : "hash",
  "fields" : [
    "a",
    "b"
  ],
  "selectivityEstimate" : 1,
  "unique" : true,
  "sparse" : false,
  "isNewlyCreated" : true,
  "error" : false,
  "code" : 201
}
```

Example:

Creating a non-unique hash index

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/index?collection=products <<EOF
{
  "type" : "hash",
  "unique" : false,
  "fields" : [
    "a",
    "b"
  ]
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8

{
  "id" : "products/10433",
  "type" : "hash",
  "fields" : [
    "a",
    "b"
  ],
  "selectivityEstimate" : 1,
  "unique" : false,
  "sparse" : false,
  "isNewlyCreated" : true,
  "error" : false,
  "code" : 201
}
```

Example:

Creating a sparse index

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/index?collection=products <<EOF
{
  "type" : "hash",
  "unique" : false,
  "sparse" : true,
  "fields" : [
    "a"
  ]
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8

{
  "id" : "products/10457",
  "type" : "hash",
  "fields" : [
    "a"
  ],
  "selectivityEstimate" : 1,
  "unique" : false,
  "sparse" : true,
  "isNewlyCreated" : true,
  "error" : false,
  "code" : 201
}
```

Return Codes

- *200*: If the index already exists, then a *HTTP 200* is returned.
- *201*: If the index does not already exist and could be created, then a *HTTP 201* is returned.
- *400*: If the collection already contains documents and you try to create a unique hash index in such a way that there are documents violating the uniqueness, then a *HTTP 400* is returned.
- *404*: If the *collection-name* is unknown, then a *HTTP 404* is returned.

Examples

Creating an unique constraint

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/index?
collection=products <<EOF
{
  "type" : "hash",
  "unique" : true,
  "fields" : [
    "a",
    "b"
  ]
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8
```

show response body

Creating a non-unique hash index

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/index?
collection=products <<EOF
{
  "type" : "hash",
  "unique" : false,
  "fields" : [
    "a",
    "b"
  ]
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8
```

show response body

Creating a sparse index

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/index?
collection=products <<EOF
{
  "type" : "hash",
  "unique" : false,
  "sparse" : true,
  "fields" : [
    "a"
  ]
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8
```

show response body

Simple query by-example

returns all documents of a collection matching a given example

`PUT /_api/simple/by-example`

A json post document with these Properties is required:

- **skip**: The number of documents to skip in the query (optional).
- **batchSize**: maximum number of result documents to be transferred from the server to the client in one roundtrip. If this attribute is not set, a server-controlled default value will be used. A *batchSize* value of *0* is disallowed.
- **limit**: The maximal amount of documents to return. The *skip* is applied before the *limit* restriction. (optional)
- **example**: The example document.
- **collection**: The name of the collection to query.

This will find all documents matching a given example.

Returns a cursor containing the result, see Http Cursor for details.

Example:

Matching an attribute

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/by-example <<EOF
{
  "collection" : "products",
  "example" : {
    "i" : 1
  }
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8

{
  "result" : [
    {
      "_key" : "10659",
      "_id" : "products/10659",
      "_rev" : "_U2FCn1W---",
      "a" : {
        "k" : 1,
        "j" : 1
      },
      "i" : 1
    },
    {
      "_key" : "10663",
      "_id" : "products/10663",
      "_rev" : "_U2FCn1a---",
      "a" : {
        "j" : 1
      },
      "i" : 1
    },
    {
      "_key" : "10666",
      "_id" : "products/10666",
      "_rev" : "_U2FCn1a--_",
      "i" : 1
    },
    {
      "_key" : "10669",
      "_id" : "products/10669",
      "_rev" : "_U2FCn1e---",
      "a" : {
        "k" : 2,
        "j" : 2
      },
      "i" : 1
    }
  ],
  "hasMore" : false,
  "count" : 4,
  "error" : false,
  "code" : 201
}
```

Example:

Matching an attribute which is a sub-document

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/by-example <<EOF
{
  "collection" : "products",
  "example" : {
    "a.j" : 1
  }
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8

{
  "result" : [
    {
      "_key" : "10677",
      "_id" : "products/10677",
      "_rev" : "_U2FCn8----",
      "a" : {
        "k" : 1,
        "j" : 1
      },
      "i" : 1
    },
    {
      "_key" : "10681",
      "_id" : "products/10681",
      "_rev" : "_U2FCn8---_",
      "a" : {
        "j" : 1
      },
      "i" : 1
    }
  ],
  "hasMore" : false,
  "count" : 2,
  "error" : false,
  "code" : 201
}
```

Example:

Matching an attribute within a sub-document

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/by-example <<EOF
{
  "collection" : "products",
  "example" : {
    "a" : {
      "j" : 1
    }
  }
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8

{
  "result" : [
    {
      "_key" : "10699",
      "_id" : "products/10699",
      "_rev" : "_U2FCoDW---",
      "a" : {
        "j" : 1
      },
      "i" : 1
    }
  ],
  "hasMore" : false,
  "count" : 1,
  "error" : false,
  "code" : 201
}
```

Return Codes

- *201*: is returned if the query was executed successfully.
- *400*: is returned if the body does not contain a valid JSON representation of a query. The response body contains an error document in this case.
- *404*: is returned if the collection specified by *collection* is unknown. The response body contains an error document in this case.

Examples

Matching an attribute

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/by-example <<EOF
{
  "collection" : "products",
  "example" : {
    "i" : 1
  }
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8
```

show response body

Matching an attribute which is a sub-document

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/by-
example <<EOF
{
  "collection" : "products",
  "example" : {
    "a.j" : 1
  }
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8
```

show response body

Matching an attribute within a sub-document

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/by-
example <<EOF
{
  "collection" : "products",
  "example" : {
    "a" : {
      "j" : 1
    }
  }
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8
```

show response body

Find documents matching an example

returns one document of a collection matching a given example

`PUT /_api/simple/first-example`

A json post document with these Properties is required:

- **example**: The example document.
- **collection**: The name of the collection to query.

This will return the first document matching a given example.

Returns a result containing the document or *HTTP 404* if no document matched the example.

If more than one document in the collection matches the specified example, only one of these documents will be returned, and it is undefined which of the matching documents is returned.

Example:

If a matching document was found

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/first-example <<EOF
{
  "collection" : "products",
  "example" : {
    "i" : 1
  }
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "document" : {
    "_key" : "10720",
    "_id" : "products/10720",
    "_rev" : "_U2FCoKu---",
    "i" : 1
  },
  "error" : false,
  "code" : 200
}
```

Example:

If no document was found

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/first-example <<EOF
{
  "collection" : "products",
  "example" : {
    "l" : 1
  }
}
EOF

HTTP/1.1 404 Not Found
content-type: application/json; charset=utf-8

{
  "error" : true,
  "code" : 404,
  "errorNum" : 404,
  "errorMessage" : "no match"
}
```

Return Codes

- *200*: is returned when the query was successfully executed.
- *400*: is returned if the body does not contain a valid JSON representation of a query. The response body contains an error document in this case.
- *404*: is returned if the collection specified by *collection* is unknown. The response body contains an error document in this case.

Examples

If a matching document was found

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/first-
example <<EOF
{
  "collection" : "products",
  "example" : {
    "i" : 1
  }
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

If no document was found

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/first-
example <<EOF
{
  "collection" : "products",
  "example" : {
    "l" : 1
  }
}
EOF

HTTP/1.1 404 Not Found
content-type: application/json; charset=utf-8
```

show response body

Working with Skiplist Indexes

If a suitable skip-list index exists, then /_api/simple/range and other operations will use this index to execute queries.

Create skip list

creates a skip-list

`POST /_api/index#skiplist`

Query Parameters

- *collection-name* (required): The collection name.

A json post document with these Properties is required:

- **fields** (string): an array of attribute paths.
- **unique**: if *true*, then create a unique index.
- **type**: must be equal to *"skiplist"*.
- **sparse**: if *true*, then create a sparse index.

Creates a skip-list index for the collection *collection-name*, if it does not already exist. The call expects an object containing the index details.

In a sparse index all documents will be excluded from the index that do not contain at least one of the specified index attributes (i.e. *fields*) or that have a value of *null* in any of the specified index attributes. Such documents will not be indexed, and not be taken into account for uniqueness checks if the *unique* flag is set.

In a non-sparse index, these documents will be indexed (for non-present indexed attributes, a value of *null* will be used) and will be taken into account for uniqueness checks if the *unique* flag is set.

Note: unique indexes on non-shard keys are not supported in a cluster.

Example:

Creating a skiplist index

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/index?collection=products <<EOF
{
  "type" : "skiplist",
  "unique" : false,
  "fields" : [
    "a",
    "b"
  ]
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8

{
  "id" : "products/10445",
  "type" : "skiplist",
  "fields" : [
    "a",
    "b"
  ],
  "unique" : false,
  "sparse" : false,
  "isNewlyCreated" : true,
  "error" : false,
  "code" : 201
}
```

Example:

Creating a sparse skiplist index

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/index?collection=products <<EOF
{
  "type" : "skiplist",
  "unique" : false,
  "sparse" : true,
  "fields" : [
    "a"
  ]
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8

{
  "id" : "products/10469",
  "type" : "skiplist",
  "fields" : [
    "a"
  ],
  "unique" : false,
  "sparse" : true,
  "isNewlyCreated" : true,
  "error" : false,
  "code" : 201
}
```

Return Codes

- *200*: If the index already exists, then a *HTTP 200* is returned.
- *201*: If the index does not already exist and could be created, then a *HTTP 201* is returned.
- *400*: If the collection already contains documents and you try to create a unique skip-list index in such a way that there are documents violating the uniqueness, then a *HTTP 400* is returned.
- *404*: If the *collection-name* is unknown, then a *HTTP 404* is returned.

Examples

Creating a skiplist index

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/index?
collection=products <<EOF
{
  "type" : "skiplist",
  "unique" : false,
  "fields" : [
    "a",
    "b"
  ]
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8
```

show response body

Creating a sparse skiplist index

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/index?
collection=products <<EOF
{
  "type" : "skiplist",
  "unique" : false,
  "sparse" : true,
  "fields" : [
    "a"
  ]
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8
```

show response body

Working with Persistent Indexes

If a suitable persistent index exists, then /_api/simple/range and other operations will use this index to execute queries.

Create a persistent index

creates a persistent index

`POST /_api/index#persistent`

Query Parameters

- *collection-name* (required): The collection name.

A json post document with these Properties is required:

- **fields** (string): an array of attribute paths.
- **unique**: if *true*, then create a unique index.
- **type**: must be equal to *"persistent"*.
- **sparse**: if *true*, then create a sparse index.

Creates a persistent index for the collection *collection-name*, if it does not already exist. The call expects an object containing the index details.

In a sparse index all documents will be excluded from the index that do not contain at least one of the specified index attributes (i.e. *fields*) or that have a value of *null* in any of the specified index attributes. Such documents will not be indexed, and not be taken into account for uniqueness checks if the *unique* flag is set.

In a non-sparse index, these documents will be indexed (for non-present indexed attributes, a value of *null* will be used) and will be taken into account for uniqueness checks if the *unique* flag is set.

Note: unique indexes on non-shard keys are not supported in a cluster.

Example:

Creating a persistent index

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/index?collection=products <<EOF
{
  "type" : "persistent",
  "unique" : false,
  "fields" : [
    "a",
    "b"
  ]
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8

{
  "id" : "products/10439",
  "type" : "persistent",
  "fields" : [
    "a",
    "b"
  ],
  "unique" : false,
  "sparse" : false,
  "isNewlyCreated" : true,
  "error" : false,
  "code" : 201
}
```

Example:

Creating a sparse persistent index

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/index?collection=products <<EOF
{
  "type" : "persistent",
  "unique" : false,
  "sparse" : true,
  "fields" : [
    "a"
  ]
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8

{
  "id" : "products/10463",
  "type" : "persistent",
  "fields" : [
    "a"
  ],
  "unique" : false,
  "sparse" : true,
  "isNewlyCreated" : true,
  "error" : false,
  "code" : 201
}
```

Return Codes

- *200*: If the index already exists, then a *HTTP 200* is returned.
- *201*: If the index does not already exist and could be created, then a *HTTP 201* is returned.
- *400*: If the collection already contains documents and you try to create a unique persistent index in such a way that there are documents violating the uniqueness, then a *HTTP 400* is returned.
- *404*: If the *collection-name* is unknown, then a *HTTP 404* is returned.

Examples

Creating a persistent index

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/index?
collection=products <<EOF
{
  "type" : "persistent",
  "unique" : false,
  "fields" : [
    "a",
    "b"
  ]
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8
```

show response body

Creating a sparse persistent index

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/index?
collection=products <<EOF
{
  "type" : "persistent",
  "unique" : false,
  "sparse" : true,
  "fields" : [
    "a"
  ]
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8
```

show response body

Working with Geo Indexes

Create geo-spatial index

creates a geo index

`POST /_api/index#geo`

Query Parameters

- *collection* (required): The collection name.

A json post document with these Properties is required:

- **fields** (string): An array with one or two attribute paths. If it is an array with one attribute path *location*, then a geo-spatial index on all documents is created using *location* as path to the coordinates. The value of the attribute must be an array with at least two double values. The array must contain the latitude (first value) and the longitude (second value). All documents, which do not have the attribute path or with value that are not suitable, are ignored. If it is an array with two attribute paths *latitude* and *longitude*, then a geo-spatial index on all documents is created using *latitude* and *longitude* as paths the latitude and the longitude. The value of the attribute *latitude* and of the attribute *longitude* must a double. All documents, which do not have the attribute paths or which values are not suitable, are ignored.
- **type**: must be equal to *"geo"*.
- **geoJson**: If a geo-spatial index on a *location* is constructed and *geoJson* is *true*, then the order within the array is longitude followed by latitude. This corresponds to the format described in http://geojson.org/geojson-spec.html#positions

NOTE Swagger examples won't work due to the anchor.

Creates a geo-spatial index in the collection *collection-name*, if it does not already exist. Expects an object containing the index details.

Geo indexes are always sparse, meaning that documents that do not contain the index attributes or have non-numeric values in the index attributes will not be indexed.

Example:

Creating a geo index with a location attribute

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/index?collection=products <<EOF
{
  "type" : "geo",
  "fields" : [
    "b"
  ]
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8

{
  "id" : "products/10421",
  "type" : "geo1",
  "fields" : [
    "b"
  ],
  "geoJson" : false,
  "constraint" : false,
  "unique" : false,
  "ignoreNull" : true,
  "sparse" : true,
  "isNewlyCreated" : true,
  "error" : false,
  "code" : 201
}
```

Example:

Creating a geo index with latitude and longitude attributes

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/index?collection=products <<EOF
{
  "type" : "geo",
  "fields" : [
    "e",
    "f"
  ]
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8

{
  "id" : "products/10415",
  "type" : "geo2",
  "fields" : [
    "e",
    "f"
  ],
  "constraint" : false,
  "unique" : false,
  "ignoreNull" : true,
  "sparse" : true,
  "isNewlyCreated" : true,
  "error" : false,
  "code" : 201
}
```

Return Codes

- *200*: If the index already exists, then a *HTTP 200* is returned.
- *201*: If the index does not already exist and could be created, then a *HTTP 201* is returned.
- *404*: If the *collection-name* is unknown, then a *HTTP 404* is returned.

Examples

Creating a geo index with a location attribute

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/index?
collection=products <<EOF
{
  "type" : "geo",
  "fields" : [
    "b"
  ]
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8
```

show response body

Creating a geo index with latitude and longitude attributes

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/index?
collection=products <<EOF
{
  "type" : "geo",
  "fields" : [
    "e",
    "f"
  ]
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8
```

show response body

Returns documents near a coordinate

returns all documents of a collection near a given location

`PUT /_api/simple/near`

A json post document with these Properties is required:

- **distance**: If given, the attribute key used to return the distance to the given coordinate. (optional). If specified, distances are returned in meters.
- **skip**: The number of documents to skip in the query. (optional)
- **longitude**: The longitude of the coordinate.
- **limit**: The maximal amount of documents to return. The *skip* is applied before the *limit* restriction. The default is 100. (optional)
- **collection**: The name of the collection to query.
- **latitude**: The latitude of the coordinate.
- **geo**: If given, the identifier of the geo-index to use. (optional)

The default will find at most 100 documents near the given coordinate. The returned array is sorted according to the distance, with the nearest document being first in the return array. If there are near documents of equal distance, documents are chosen randomly from this set until the limit is reached.

In order to use the *near* operator, a geo index must be defined for the collection. This index also defines which attribute holds the coordinates for the document. If you have more than one geo-spatial index, you can use the *geo* field to select a particular index.

Returns a cursor containing the result, see Http Cursor for details.

Note: the *near* simple query is **deprecated** as of ArangoDB 2.6. This API may be removed in future versions of ArangoDB. The preferred way for retrieving documents from a collection using the near operator is to issue an AQL query using the *NEAR* function as follows:

```
FOR doc IN NEAR(@@collection, @latitude, @longitude, @limit)
  RETURN doc`
```

Example:

Without distance

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/near <<EOF
{
  "collection" : "products",
  "latitude" : 0,
  "longitude" : 0,
  "skip" : 1,
  "limit" : 2
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8

{
  "result" : [
    {
      "_key" : "10847",
      "_id" : "products/10847",
      "_rev" : "_U2FCpKK--B",
      "name" : "Name/0.002/",
      "loc" : [
        0.002,
        0
      ]
    },
    {
      "_key" : "10841",
      "_id" : "products/10841",
      "_rev" : "_U2FCpKK--_",
      "name" : "Name/-0.002/",
      "loc" : [
        -0.002,
        0
      ]
    }
  ],
  "hasMore" : false,
  "count" : 2,
  "error" : false,
  "code" : 201
}
```

Example:

With distance

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/near <<EOF
{
  "collection" : "products",
  "latitude" : 0,
  "longitude" : 0,
  "skip" : 1,
  "limit" : 3,
  "distance" : "distance"
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8

{
  "result" : [
    {
      "distance" : 222.38985328911744,
      "_id" : "products/10884",
      "_key" : "10884",
      "_rev" : "_U2FCpXW---",
      "loc" : [
        -0.002,
        0
      ],
      "name" : "Name/-0.002/"
    },
    {
      "distance" : 222.38985328911744,
      "_id" : "products/10890",
      "_key" : "10890",
      "_rev" : "_U2FCpXW--A",
      "loc" : [
        0.002,
        0
      ],
      "name" : "Name/0.002/"
    },
    {
      "distance" : 444.779706578235,
      "_id" : "products/10881",
      "_key" : "10881",
      "_rev" : "_U2FCpXS--_",
      "loc" : [
        -0.004,
        0
      ],
      "name" : "Name/-0.004/"
    }
  ],
  "hasMore" : false,
  "count" : 3,
  "error" : false,
  "code" : 201
}
```

Return Codes

- *201*: is returned if the query was executed successfully.
- *400*: is returned if the body does not contain a valid JSON representation of a query. The response body contains an error document in this case.
- *404*: is returned if the collection specified by *collection* is unknown. The response body contains an error document in this case.

Examples

Without distance

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/near
<<EOF
{
  "collection" : "products",
  "latitude" : 0,
  "longitude" : 0,
  "skip" : 1,
  "limit" : 2
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8
```

show response body

With distance

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/near
<<EOF
{
  "collection" : "products",
  "latitude" : 0,
  "longitude" : 0,
  "skip" : 1,
  "limit" : 3,
  "distance" : "distance"
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8
```

show response body

Find documents within a radius around a coordinate

returns all documents of a collection within a given radius

`PUT /_api/simple/within`

A json post document with these Properties is required:

- **distance**: If given, the attribute key used to return the distance to the given coordinate. (optional). If specified, distances are returned in meters.
- **skip**: The number of documents to skip in the query. (optional)
- **longitude**: The longitude of the coordinate.
- **radius**: The maximal radius (in meters).
- **collection**: The name of the collection to query.
- **latitude**: The latitude of the coordinate.
- **limit**: The maximal amount of documents to return. The *skip* is applied before the *limit* restriction. The default is 100. (optional)
- **geo**: If given, the identifier of the geo-index to use. (optional)

This will find all documents within a given radius around the coordinate (*latitude*, *longitude*). The returned list is sorted by distance.

In order to use the *within* operator, a geo index must be defined for the collection. This index also defines which attribute holds the coordinates for the document. If you have more than one geo-spatial index, you can use the *geo* field to select a particular index.

Returns a cursor containing the result, see Http Cursor for details.

Note: the *within* simple query is **deprecated** as of ArangoDB 2.6. This API may be removed in future versions of ArangoDB. The preferred way for retrieving documents from a collection using the near operator is to issue an AQL query using the *WITHIN* function as follows:

```
FOR doc IN WITHIN(@@collection, @latitude, @longitude, @radius, @distanceAttributeName)
  RETURN doc
```

Example:

Without distance

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/near <<EOF
{
  "collection" : "products",
  "latitude" : 0,
  "longitude" : 0,
  "skip" : 1,
  "limit" : 2,
  "radius" : 500
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8

{
  "result" : [
    {
      "_key" : "11164",
      "_id" : "products/11164",
      "_rev" : "_U2FCrGO--_",
      "name" : "Name/0.002/",
      "loc" : [
        0.002,
        0
      ]
    },
    {
      "_key" : "11158",
      "_id" : "products/11158",
      "_rev" : "_U2FCrGK--_",
      "name" : "Name/-0.002/",
      "loc" : [
        -0.002,
        0
      ]
    }
  ],
  "hasMore" : false,
  "count" : 2,
  "error" : false,
  "code" : 201
}
```

Example:

With distance

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/near <<EOF
{
  "collection" : "products",
  "latitude" : 0,
  "longitude" : 0,
  "skip" : 1,
  "limit" : 3,
  "distance" : "distance",
  "radius" : 300
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8

{
  "result" : [
    {
      "distance" : 222.38985328911744,
      "_id" : "products/11201",
      "_key" : "11201",
      "_rev" : "_U2FCrSu--A",
      "loc" : [
        -0.002,
        0
      ],
      "name" : "Name/-0.002/"
    },
    {
      "distance" : 222.38985328911744,
      "_id" : "products/11207",
      "_key" : "11207",
      "_rev" : "_U2FCrSy--_",
      "loc" : [
        0.002,
        0
      ],
      "name" : "Name/0.002/"
    },
    {
      "distance" : 444.779706578235,
      "_id" : "products/11198",
      "_key" : "11198",
      "_rev" : "_U2FCrSu--_",
      "loc" : [
        -0.004,
        0
      ],
      "name" : "Name/-0.004/"
    }
  ],
  "hasMore" : false,
  "count" : 3,
  "error" : false,
  "code" : 201
}
```

Return Codes

- *201*: is returned if the query was executed successfully.
- *400*: is returned if the body does not contain a valid JSON representation of a query. The response body contains an error document in this case.
- *404*: is returned if the collection specified by *collection* is unknown. The response body contains an error document in this case.

Examples

Without distance

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/near
<<EOF
{
  "collection" : "products",
  "latitude" : 0,
  "longitude" : 0,
  "skip" : 1,
  "limit" : 2,
  "radius" : 500
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8
```

show response body

With distance

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/near
<<EOF
{
  "collection" : "products",
  "latitude" : 0,
  "longitude" : 0,
  "skip" : 1,
  "limit" : 3,
  "distance" : "distance",
  "radius" : 300
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8
```

show response body

Fulltext

If a fulltext index exists, then /_api/simple/fulltext will use this index to execute the specified fulltext query.

Create fulltext index

creates a fulltext index

`POST /_api/index#fulltext`

Query Parameters

- *collection-name* (required): The collection name.

A json post document with these Properties is required:

- **fields** (string): an array of attribute names. Currently, the array is limited to exactly one attribute.
- **type**: must be equal to *"fulltext"*.
- **minLength**: Minimum character length of words to index. Will default to a server-defined value if unspecified. It is thus recommended to set this value explicitly when creating the index.

NOTE Swagger examples won't work due to the anchor.

Creates a fulltext index for the collection *collection-name*, if it does not already exist. The call expects an object containing the index details.

Example:

Creating a fulltext index

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/index?collection=products <<EOF
{
  "type" : "fulltext",
  "fields" : [
    "text"
  ]
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8

{
  "id" : "products/10427",
  "type" : "fulltext",
  "fields" : [
    "text"
  ],
  "unique" : false,
  "sparse" : true,
  "minLength" : 2,
  "isNewlyCreated" : true,
  "error" : false,
  "code" : 201
}
```

Return Codes

- *200*: If the index already exists, then a *HTTP 200* is returned.
- *201*: If the index does not already exist and could be created, then a *HTTP 201* is returned.
- *404*: If the *collection-name* is unknown, then a *HTTP 404* is returned.

Examples

Creating a fulltext index

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/index?
collection=products <<EOF
{
  "type" : "fulltext",
  "fields" : [
    "text"
  ]
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8
```

show response body

Fulltext index query

returns documents of a collection as a result of a fulltext query

`PUT /_api/simple/fulltext`

A json post document with these Properties is required:

- **index**: The identifier of the fulltext-index to use.
- **attribute**: The attribute that contains the texts.
- **collection**: The name of the collection to query.
- **limit**: The maximal amount of documents to return. The *skip* is applied before the *limit* restriction. (optional)
- **skip**: The number of documents to skip in the query (optional).
- **query**: The fulltext query. Please refer to Fulltext queries for details.

This will find all documents from the collection that match the fulltext query specified in *query*.

In order to use the *fulltext* operator, a fulltext index must be defined for the collection and the specified attribute.

Returns a cursor containing the result, see Http Cursor for details.

Note: the *fulltext* simple query is **deprecated** as of ArangoDB 2.6. This API may be removed in future versions of ArangoDB. The preferred way for retrieving documents from a collection using the near operator is to issue an AQL query using the *FULLTEXT* AQL function as follows:

```
FOR doc IN FULLTEXT(@@collection, @attributeName, @queryString, @limit)
  RETURN doc
```

Example:

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/fulltext <<EOF
{
  "collection" : "products",
  "attribute" : "text",
  "query" : "word"
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8

{
  "result" : [
    {
      "_key" : "10752",
      "_id" : "products/10752",
      "_rev" : "_U2FCola--_",
      "text" : "this text also has a word"
    },
    {
      "_key" : "10748",
      "_id" : "products/10748",
      "_rev" : "_U2FCola---",
      "text" : "this text contains word"
    }
  ],
  "hasMore" : false,
  "count" : 2,
  "error" : false,
  "code" : 201
}
```

Return Codes

- *201*: is returned if the query was executed successfully.
- *400*: is returned if the body does not contain a valid JSON representation of a query. The response body contains an error document in this case.
- *404*: is returned if the collection specified by *collection* is unknown. The response body contains an error document in this case.

Examples

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/simple/fulltext
<<EOF
{
  "collection" : "products",
  "attribute" : "text",
  "query" : "word"
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8
```

show response body

HTTP Interface for Transactions

Transactions

ArangoDB's transactions are executed on the server. Transactions can be initiated by clients by sending the transaction description for execution to the server.

Transactions in ArangoDB do not offer separate *BEGIN*, *COMMIT* and *ROLLBACK* operations as they are available in many other database products. Instead, ArangoDB transactions are described by a JavaScript function, and the code inside the JavaScript function will then be executed transactionally. At the end of the function, the transaction is automatically committed, and all changes done by the transaction will be persisted. If an exception is thrown during transaction execution, all operations performed in the transaction are rolled back.

For a more detailed description of how transactions work in ArangoDB please refer to Transactions.

Execute transaction

execute a server-side transaction

`POST /_api/transaction`

A json post document with these Properties is required:

- **action**: the actual transaction operations to be executed, in the form of stringified JavaScript code. The code will be executed on server side, with late binding. It is thus critical that the code specified in *action* properly sets up all the variables it needs. If the code specified in *action* ends with a return statement, the value returned will also be returned by the REST API in the *result* attribute if the transaction committed successfully.
- **params**: optional arguments passed to *action*.
- **collections**: *collections* must be a JSON object that can have either or both sub-attributes *read* and *write*, each being an array of collection names or a single collection name as string. Collections that will be written to in the transaction must be declared with the *write* attribute or it will fail, whereas non-declared collections from which is solely read will be added lazily. The optional sub-attribute *allowImplicit* can be set to *false* to let transactions fail in case of undeclared collections for reading. Collections for reading should be fully declared if possible, to avoid deadlocks.
- **lockTimeout**: an optional numeric value that can be used to set a timeout for waiting on collection locks. If not specified, a default value will be used. Setting *lockTimeout* to *0* will make ArangoDB not time out waiting for a lock.
- **waitForSync**: an optional boolean flag that, if set, will force the transaction to write all data to disk before returning.

Contains the *collections* and *action*.

The transaction description must be passed in the body of the POST request.

If the transaction is fully executed and committed on the server, *HTTP 200* will be returned. Additionally, the return value of the code defined in *action* will be returned in the *result* attribute.

For successfully committed transactions, the returned JSON object has the following properties:

- *error*: boolean flag to indicate if an error occurred (*false* in this case)
- *code*: the HTTP status code
- *result*: the return value of the transaction

If the transaction specification is either missing or malformed, the server will respond with *HTTP 400*.

The body of the response will then contain a JSON object with additional error details. The object has the following attributes:

- *error*: boolean flag to indicate that an error occurred (*true* in this case)
- *code*: the HTTP status code
- *errorNum*: the server error number
- *errorMessage*: a descriptive error message

If a transaction fails to commit, either by an exception thrown in the *action* code, or by an internal error, the server will respond with an error. Any other errors will be returned with any of the return codes *HTTP 400*, *HTTP 409*, or *HTTP 500*.

Example:

Executing a transaction on a single collection

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/transaction <<EOF
{
  "collections" : {
    "write" : "products"
  },
  "action" : "function () { var db = require('@arangodb').db; db.products.save({});  return db.products.count(); }"
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "result" : 1,
  "error" : false,
  "code" : 200
}
```

Example:

Executing a transaction using multiple collections

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/transaction <<EOF
{
  "collections" : {
    "write" : [
      "products",
      "materials"
    ]
  },
  "action" : "function () {var db = require('@arangodb').db;db.products.save({});db.materials.save({});return 'worked
!';}"
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "result" : "worked!",
  "error" : false,
  "code" : 200
}
```

Example:

Aborting a transaction due to an internal error

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/transaction <<EOF
{
  "collections" : {
    "write" : "products"
  },
  "action" : "function () {var db = require('@arangodb').db;db.products.save({ _key: 'abc'});db.products.save({ _key:
 'abc'});}"
}
EOF

HTTP/1.1 400 Bad Request
content-type: application/json; charset=utf-8

{
  "exception" : "ArangoError 1210: unique constraint violated",
  "stacktrace" : [
    "ArangoError: unique constraint violated",
    "    at Error (native)",
    "    at eval (eval at post_api_transaction (/home/jenkins/stable-3.1/js/actions/api-transaction.js:48:16), <anony
mous>:3:100)",
    "    at eval (eval at post_api_transaction (/home/jenkins/stable-3.1/js/actions/api-transaction.js:48:16), <anony
mous>:3:123)",
    "    at post_api_transaction (/home/jenkins/stable-3.1/js/actions/api-transaction.js:48:16)",
    "    at Function.actions.defineHttp.callback (/home/jenkins/stable-3.1/js/actions/api-transaction.js:64:11)"
  ],
  "message" : "unique constraint violated",
  "error" : true,
  "code" : 400,
  "errorNum" : 1210,
  "errorMessage" : "unique constraint violated"
}
```

Example:

Aborting a transaction by explicitly throwing an exception

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/transaction <<EOF
{
  "collections" : {
    "read" : "products"
  },
  "action" : "function () { throw 'doh!'; }"
}
EOF

HTTP/1.1 500 Internal Server Error
content-type: application/json; charset=utf-8

{
  "exception" : "doh!",
  "error" : true,
  "code" : 500,
  "errorNum" : 500,
  "errorMessage" : "internal server error"
}
```

Example:

Referring to a non-existing collection

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/transaction <<EOF
{
  "collections" : {
    "read" : "products"
  },
  "action" : "function () { return true; }"
}
EOF

HTTP/1.1 404 Not Found
content-type: application/json; charset=utf-8

{
  "exception" : "ArangoError 1203: collection not found",
  "stacktrace" : [
    "ArangoError: collection not found",
    "    at Error (native)",
    "    at post_api_transaction (/home/jenkins/stable-3.1/js/actions/api-transaction.js:48:16)",
    "    at Function.actions.defineHttp.callback (/home/jenkins/stable-3.1/js/actions/api-transaction.js:64:11)"
  ],
  "message" : "collection not found",
  "error" : true,
  "code" : 404,
  "errorNum" : 1203,
  "errorMessage" : "collection not found"
}
```

Return Codes

- *200*: If the transaction is fully executed and committed on the server, *HTTP 200* will be returned.
- *400*: If the transaction specification is either missing or malformed, the server will respond with *HTTP 400*.
- *404*: If the transaction specification contains an unknown collection, the server will respond with *HTTP 404*.
- *500*: Exceptions thrown by users will make the server respond with a return code of *HTTP 500*

Examples

Executing a transaction on a single collection

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/transaction
<<EOF
{
  "collections" : {
    "write" : "products"
  },
  "action" : "function () { var db = require('@arangodb').db; db.products.save({});
return db.products.count(); }"
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Executing a transaction using multiple collections

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/transaction
<<EOF
{
  "collections" : {
    "write" : [
      "products",
      "materials"
    ]
  },
  "action" : "function () {var db =
require('@arangodb').db;db.products.save({});db.materials.save({});return 'worked!';}"
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Aborting a transaction due to an internal error

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/transaction
<<EOF
{
  "collections" : {
    "write" : "products"
  },
  "action" : "function () {var db = require('@arangodb').db;db.products.save({ _key:
'abc'});db.products.save({ _key: 'abc'});}"
}
EOF

HTTP/1.1 400 Bad Request
content-type: application/json; charset=utf-8
```

show response body

Aborting a transaction by explicitly throwing an exception

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/transaction
<<EOF
{
  "collections" : {
    "read" : "products"
  },
  "action" : "function () { throw 'doh!'; }"
}
EOF

HTTP/1.1 500 Internal Server Error
content-type: application/json; charset=utf-8
```

show response body

Referring to a non-existing collection

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/transaction
<<EOF
{
  "collections" : {
    "read" : "products"
  },
  "action" : "function () { return true; }"
}
EOF

HTTP/1.1 404 Not Found
content-type: application/json; charset=utf-8
```

show response body

HTTP Interface for Replication

Replication

This is an introduction to ArangoDB's HTTP replication interface. The replication architecture and components are described in more details in Replication.

The HTTP replication interface serves four main purposes:

- fetch initial data from a server (e.g. for a backup, or for the initial synchronization of data before starting the continuous replication applier)
- querying the state of a master
- fetch continuous changes from a master (used for incremental synchronization of changes)
- administer the replication applier (starting, stopping, configuring, querying state) on a slave

Please note that all replication operations work on a per-database level. If an ArangoDB server contains more than one database, the replication system must be configured individually per database, and replicating the data of multiple databases will require multiple operations.

Replication Dump Commands

The *inventory* method can be used to query an ArangoDB database's current set of collections plus their indexes. Clients can use this method to get an overview of which collections are present in the database. They can use this information to either start a full or a partial synchronization of data, e.g. to initiate a backup or the incremental data synchronization.

Return inventory of collections and indexes

Returns an overview of collections and their indexes

`GET /_api/replication/inventory`

Query Parameters

- *includeSystem* (optional): Include system collections in the result. The default value is *true*.

Returns the array of collections and indexes available on the server. This array can be used by replication clients to initiate an initial sync with the server.

The response will contain a JSON object with the *collection* and *state* and *tick* attributes.

collections is a array of collections with the following sub-attributes:

- *parameters*: the collection properties
- *indexes*: a array of the indexes of a the collection. Primary indexes and edges indexes are not included in this array.

The *state* attribute contains the current state of the replication logger. It contains the following sub-attributes:

- *running*: whether or not the replication logger is currently active. Note: since ArangoDB 2.2, the value will always be *true*
- *lastLogTick*: the value of the last tick the replication logger has written
- *time*: the current time on the server

Replication clients should note the *lastLogTick* value returned. They can then fetch collections' data using the dump method up to the value of lastLogTick, and query the continuous replication log for log events after this tick value.

To create a full copy of the collections on the server, a replication client can execute these steps:

- call the */inventory* API method. This returns the *lastLogTick* value and the array of collections and indexes from the server.
- for each collection returned by */inventory*, create the collection locally and call */dump* to stream the collection data to the client, up to the value of *lastLogTick*. After that, the client can create the indexes on the collections as they were reported by */inventory*.

If the clients wants to continuously stream replication log events from the logger server, the following additional steps need to be carried out:

- the client should call */logger-follow* initially to fetch the first batch of replication events that were logged after the client's call to */inventory*.

  The call to */logger-follow* should use a *from* parameter with the value of the *lastLogTick* as reported by */inventory*. The call to */logger-follow* will return the *x-arango-replication-lastincluded* which will contain the last tick value included in the response.

- the client can then continuously call */logger-follow* to incrementally fetch new replication events that occurred after the last transfer.

  Calls should use a *from* parameter with the value of the *x-arango-replication-lastincluded* header of the previous response. If there are no more replication events, the response will be empty and clients can go to sleep for a while and try again later.

Note: on a coordinator, this request must have the query parameter *DBserver* which must be an ID of a DBserver. The very same request is forwarded synchronously to that DBserver. It is an error if this attribute is not bound in the coordinator case.

Example:

```
shell> curl --dump - http://localhost:8529/_api/replication/inventory

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "collections" : [
    {
      "parameters" : {
        "allowUserKeys" : true,
        "cid" : "36",
        "count" : 0,
        "deleted" : false,
        "distributeShardsLike" : "_graphs",
        "doCompact" : true,
        "id" : "36",
        "indexBuckets" : 8,
        "indexes" : [
          {
            "fields" : [
              "_key"
            ],
            "id" : "0",
            "selectivityEstimate" : 1,
            "sparse" : false,
            "type" : "primary",
            "unique" : true
          }
        ],
        "isSystem" : true,
        "isVolatile" : false,
        "journalSize" : 1048576,
        "keyOptions" : {
          "allowUserKeys" : true,
          "type" : "traditional"
        },
        "name" : "_aqlfunctions",
        "numberOfShards" : 1,
        "planId" : "36",
        "replicationFactor" : 2,
        "shardKeys" : [
          "_key"
        ],
        "shards" : {
        },
        "status" : 0,
        "type" : 2,
        "version" : 6,
        "waitForSync" : false
      },
      "indexes" : [ ]
    },
    {
      "parameters" : {
        "allowUserKeys" : true,
        "cid" : "54",
        "count" : 0,
        "deleted" : false,
        "distributeShardsLike" : "_graphs",
        "doCompact" : true,
        "id" : "54",
        "indexBuckets" : 8,
        "indexes" : [
          {
            "fields" : [
              "_key"
            ],
            "id" : "0",
            "selectivityEstimate" : 1,
            "sparse" : false,
            "type" : "primary",
            "unique" : true
          }
        ],
        "isSystem" : true,
```

```
        "isVolatile" : false,
        "journalSize" : 1048576,
        "keyOptions" : {
          "allowUserKeys" : true,
          "type" : "traditional"
        },
        "name" : "_frontend",
        "numberOfShards" : 1,
        "planId" : "54",
        "replicationFactor" : 2,
        "shardKeys" : [
          "_key"
        ],
        "shards" : {
        },
        "status" : 0,
        "type" : 2,
        "version" : 6,
        "waitForSync" : false
      },
      "indexes" : [ ]
    },
    {
      "parameters" : {
        "allowUserKeys" : true,
        "cid" : "3",
        "count" : 0,
        "deleted" : false,
        "doCompact" : true,
        "id" : "3",
        "indexBuckets" : 8,
        "indexes" : [
          {
            "fields" : [
              "_key"
            ],
            "id" : "0",
            "selectivityEstimate" : 1,
            "sparse" : false,
            "type" : "primary",
            "unique" : true
          }
        ],
        "isSystem" : true,
        "isVolatile" : false,
        "journalSize" : 1048576,
        "keyOptions" : {
          "allowUserKeys" : true,
          "type" : "traditional"
        },
        "name" : "_graphs",
        "numberOfShards" : 1,
        "planId" : "3",
        "replicationFactor" : 2,
        "shardKeys" : [
          "_key"
        ],
        "shards" : {
        },
        "status" : 0,
        "type" : 2,
        "version" : 6,
        "waitForSync" : false
      },
      "indexes" : [ ]
    },
    {
      "parameters" : {
        "allowUserKeys" : true,
        "cid" : "20",
        "count" : 0,
        "deleted" : false,
        "distributeShardsLike" : "_graphs",
        "doCompact" : true,
        "id" : "20",
        "indexBuckets" : 8,
```

```
        "indexes" : [
          {
            "fields" : [
              "_key"
            ],
            "id" : "0",
            "selectivityEstimate" : 1,
            "sparse" : false,
            "type" : "primary",
            "unique" : true
          }
        ],
        "isSystem" : true,
        "isVolatile" : false,
        "journalSize" : 1048576,
        "keyOptions" : {
          "allowUserKeys" : true,
          "type" : "traditional"
        },
        "name" : "_modules",
        "numberOfShards" : 1,
        "planId" : "20",
        "replicationFactor" : 2,
        "shardKeys" : [
          "_key"
        ],
        "shards" : {
        },
        "status" : 0,
        "type" : 2,
        "version" : 6,
        "waitForSync" : false
      },
      "indexes" : [ ]
    },
    {
      "parameters" : {
        "allowUserKeys" : true,
        "cid" : "22",
        "count" : 0,
        "deleted" : false,
        "distributeShardsLike" : "_graphs",
        "doCompact" : true,
        "id" : "22",
        "indexBuckets" : 8,
        "indexes" : [
          {
            "fields" : [
              "_key"
            ],
            "id" : "0",
            "selectivityEstimate" : 1,
            "sparse" : false,
            "type" : "primary",
            "unique" : true
          }
        ],
        "isSystem" : true,
        "isVolatile" : false,
        "journalSize" : 4194304,
        "keyOptions" : {
          "allowUserKeys" : true,
          "type" : "traditional"
        },
        "name" : "_routing",
        "numberOfShards" : 1,
        "planId" : "22",
        "replicationFactor" : 2,
        "shardKeys" : [
          "_key"
        ],
        "shards" : {
        },
        "status" : 0,
        "type" : 2,
        "version" : 6,
```

```
        "waitForSync" : false
      },
      "indexes" : [ ]
    },
    {
      "parameters" : {
        "allowUserKeys" : true,
        "cid" : "6",
        "count" : 0,
        "deleted" : false,
        "distributeShardsLike" : "_graphs",
        "doCompact" : true,
        "id" : "6",
        "indexBuckets" : 8,
        "indexes" : [
          {
            "fields" : [
              "_key"
            ],
            "id" : "0",
            "selectivityEstimate" : 1,
            "sparse" : false,
            "type" : "primary",
            "unique" : true
          },
          {
            "fields" : [
              "user"
            ],
            "id" : "9",
            "selectivityEstimate" : 1,
            "sparse" : true,
            "type" : "hash",
            "unique" : true
          }
        ],
        "isSystem" : true,
        "isVolatile" : false,
        "journalSize" : 4194304,
        "keyOptions" : {
          "allowUserKeys" : true,
          "type" : "traditional"
        },
        "name" : "_users",
        "numberOfShards" : 1,
        "planId" : "6",
        "replicationFactor" : 2,
        "shardKeys" : [
          "user"
        ],
        "shards" : {
        },
        "status" : 3,
        "type" : 2,
        "version" : 6,
        "waitForSync" : false
      },
      "indexes" : [
        {
          "fields" : [
            "user"
          ],
          "id" : "9",
          "selectivityEstimate" : 1,
          "sparse" : true,
          "type" : "hash",
          "unique" : true
        }
      ]
    },
    {
      "parameters" : {
        "allowUserKeys" : true,
        "cid" : "98",
        "count" : 0,
        "deleted" : false,
```

```
        "doCompact" : true,
        "id" : "98",
        "indexBuckets" : 8,
        "indexes" : [
          {
            "fields" : [
              "_key"
            ],
            "id" : "0",
            "selectivityEstimate" : 1,
            "sparse" : false,
            "type" : "primary",
            "unique" : true
          }
        ],
        "isSystem" : false,
        "isVolatile" : false,
        "journalSize" : 33554432,
        "keyOptions" : {
          "allowUserKeys" : true,
          "type" : "traditional"
        },
        "name" : "animals",
        "numberOfShards" : 1,
        "planId" : "98",
        "replicationFactor" : 1,
        "shardKeys" : [
          "_key"
        ],
        "shards" : {
        },
        "status" : 0,
        "type" : 2,
        "version" : 6,
        "waitForSync" : false
      },
      "indexes" : [ ]
    },
    {
      "parameters" : {
        "allowUserKeys" : true,
        "cid" : "93",
        "count" : 0,
        "deleted" : false,
        "doCompact" : true,
        "id" : "93",
        "indexBuckets" : 8,
        "indexes" : [
          {
            "fields" : [
              "_key"
            ],
            "id" : "0",
            "selectivityEstimate" : 1,
            "sparse" : false,
            "type" : "primary",
            "unique" : true
          }
        ],
        "isSystem" : false,
        "isVolatile" : false,
        "journalSize" : 33554432,
        "keyOptions" : {
          "allowUserKeys" : true,
          "type" : "traditional"
        },
        "name" : "demo",
        "numberOfShards" : 1,
        "planId" : "93",
        "replicationFactor" : 1,
        "shardKeys" : [
          "_key"
        ],
        "shards" : {
        },
        "status" : 0,
```

```
        "type" : 2,
        "version" : 6,
        "waitForSync" : false
      },
      "indexes" : [ ]
    }
  ],
  "state" : {
    "running" : true,
    "lastLogTick" : "10547",
    "lastUncommittedLogTick" : "10549",
    "totalEvents" : 3805,
    "time" : "2017-04-18T15:53:16Z"
  },
  "tick" : "10552"
}
```

Example:

With some additional indexes:

```
shell> curl --dump - http://localhost:8529/_api/replication/inventory

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "collections" : [
    {
      "parameters" : {
        "allowUserKeys" : true,
        "cid" : "36",
        "count" : 0,
        "deleted" : false,
        "distributeShardsLike" : "_graphs",
        "doCompact" : true,
        "id" : "36",
        "indexBuckets" : 8,
        "indexes" : [
          {
            "fields" : [
              "_key"
            ],
            "id" : "0",
            "selectivityEstimate" : 1,
            "sparse" : false,
            "type" : "primary",
            "unique" : true
          }
        ],
        "isSystem" : true,
        "isVolatile" : false,
        "journalSize" : 1048576,
        "keyOptions" : {
          "allowUserKeys" : true,
          "type" : "traditional"
        },
        "name" : "_aqlfunctions",
        "numberOfShards" : 1,
        "planId" : "36",
        "replicationFactor" : 2,
        "shardKeys" : [
          "_key"
        ],
        "shards" : {
        },
        "status" : 0,
        "type" : 2,
        "version" : 6,
        "waitForSync" : false
      },
      "indexes" : [ ]
```

```
    },
    {
      "parameters" : {
        "allowUserKeys" : true,
        "cid" : "54",
        "count" : 0,
        "deleted" : false,
        "distributeShardsLike" : "_graphs",
        "doCompact" : true,
        "id" : "54",
        "indexBuckets" : 8,
        "indexes" : [
          {
            "fields" : [
              "_key"
            ],
            "id" : "0",
            "selectivityEstimate" : 1,
            "sparse" : false,
            "type" : "primary",
            "unique" : true
          }
        ],
        "isSystem" : true,
        "isVolatile" : false,
        "journalSize" : 1048576,
        "keyOptions" : {
          "allowUserKeys" : true,
          "type" : "traditional"
        },
        "name" : "_frontend",
        "numberOfShards" : 1,
        "planId" : "54",
        "replicationFactor" : 2,
        "shardKeys" : [
          "_key"
        ],
        "shards" : {
        },
        "status" : 0,
        "type" : 2,
        "version" : 6,
        "waitForSync" : false
      },
      "indexes" : [ ]
    },
    {
      "parameters" : {
        "allowUserKeys" : true,
        "cid" : "3",
        "count" : 0,
        "deleted" : false,
        "doCompact" : true,
        "id" : "3",
        "indexBuckets" : 8,
        "indexes" : [
          {
            "fields" : [
              "_key"
            ],
            "id" : "0",
            "selectivityEstimate" : 1,
            "sparse" : false,
            "type" : "primary",
            "unique" : true
          }
        ],
        "isSystem" : true,
        "isVolatile" : false,
        "journalSize" : 1048576,
        "keyOptions" : {
          "allowUserKeys" : true,
          "type" : "traditional"
        },
        "name" : "_graphs",
        "numberOfShards" : 1,
```

```
        "planId" : "3",
        "replicationFactor" : 2,
        "shardKeys" : [
          "_key"
        ],
        "shards" : {
        },
        "status" : 0,
        "type" : 2,
        "version" : 6,
        "waitForSync" : false
      },
      "indexes" : [ ]
    },
    {
      "parameters" : {
        "allowUserKeys" : true,
        "cid" : "20",
        "count" : 0,
        "deleted" : false,
        "distributeShardsLike" : "_graphs",
        "doCompact" : true,
        "id" : "20",
        "indexBuckets" : 8,
        "indexes" : [
          {
            "fields" : [
              "_key"
            ],
            "id" : "0",
            "selectivityEstimate" : 1,
            "sparse" : false,
            "type" : "primary",
            "unique" : true
          }
        ],
        "isSystem" : true,
        "isVolatile" : false,
        "journalSize" : 1048576,
        "keyOptions" : {
          "allowUserKeys" : true,
          "type" : "traditional"
        },
        "name" : "_modules",
        "numberOfShards" : 1,
        "planId" : "20",
        "replicationFactor" : 2,
        "shardKeys" : [
          "_key"
        ],
        "shards" : {
        },
        "status" : 0,
        "type" : 2,
        "version" : 6,
        "waitForSync" : false
      },
      "indexes" : [ ]
    },
    {
      "parameters" : {
        "allowUserKeys" : true,
        "cid" : "22",
        "count" : 0,
        "deleted" : false,
        "distributeShardsLike" : "_graphs",
        "doCompact" : true,
        "id" : "22",
        "indexBuckets" : 8,
        "indexes" : [
          {
            "fields" : [
              "_key"
            ],
            "id" : "0",
            "selectivityEstimate" : 1,
```

```
            "sparse" : false,
            "type" : "primary",
            "unique" : true
          }
        ],
        "isSystem" : true,
        "isVolatile" : false,
        "journalSize" : 4194304,
        "keyOptions" : {
          "allowUserKeys" : true,
          "type" : "traditional"
        },
        "name" : "_routing",
        "numberOfShards" : 1,
        "planId" : "22",
        "replicationFactor" : 2,
        "shardKeys" : [
          "_key"
        ],
        "shards" : {
        },
        "status" : 0,
        "type" : 2,
        "version" : 6,
        "waitForSync" : false
      },
      "indexes" : [ ]
    },
    {
      "parameters" : {
        "allowUserKeys" : true,
        "cid" : "6",
        "count" : 0,
        "deleted" : false,
        "distributeShardsLike" : "_graphs",
        "doCompact" : true,
        "id" : "6",
        "indexBuckets" : 8,
        "indexes" : [
          {
            "fields" : [
              "_key"
            ],
            "id" : "0",
            "selectivityEstimate" : 1,
            "sparse" : false,
            "type" : "primary",
            "unique" : true
          },
          {
            "fields" : [
              "user"
            ],
            "id" : "9",
            "selectivityEstimate" : 1,
            "sparse" : true,
            "type" : "hash",
            "unique" : true
          }
        ],
        "isSystem" : true,
        "isVolatile" : false,
        "journalSize" : 4194304,
        "keyOptions" : {
          "allowUserKeys" : true,
          "type" : "traditional"
        },
        "name" : "_users",
        "numberOfShards" : 1,
        "planId" : "6",
        "replicationFactor" : 2,
        "shardKeys" : [
          "user"
        ],
        "shards" : {
        },
```

```
        "status" : 3,
        "type" : 2,
        "version" : 6,
        "waitForSync" : false
      },
      "indexes" : [
        {
          "fields" : [
            "user"
          ],
          "id" : "9",
          "selectivityEstimate" : 1,
          "sparse" : true,
          "type" : "hash",
          "unique" : true
        }
      ]
    },
    {
      "parameters" : {
        "allowUserKeys" : true,
        "cid" : "98",
        "count" : 0,
        "deleted" : false,
        "doCompact" : true,
        "id" : "98",
        "indexBuckets" : 8,
        "indexes" : [
          {
            "fields" : [
              "_key"
            ],
            "id" : "0",
            "selectivityEstimate" : 1,
            "sparse" : false,
            "type" : "primary",
            "unique" : true
          }
        ],
        "isSystem" : false,
        "isVolatile" : false,
        "journalSize" : 33554432,
        "keyOptions" : {
          "allowUserKeys" : true,
          "type" : "traditional"
        },
        "name" : "animals",
        "numberOfShards" : 1,
        "planId" : "98",
        "replicationFactor" : 1,
        "shardKeys" : [
          "_key"
        ],
        "shards" : {
        },
        "status" : 0,
        "type" : 2,
        "version" : 6,
        "waitForSync" : false
      },
      "indexes" : [ ]
    },
    {
      "parameters" : {
        "allowUserKeys" : true,
        "cid" : "93",
        "count" : 0,
        "deleted" : false,
        "doCompact" : true,
        "id" : "93",
        "indexBuckets" : 8,
        "indexes" : [
          {
            "fields" : [
              "_key"
            ],
```

```
            "id" : "0",
            "selectivityEstimate" : 1,
            "sparse" : false,
            "type" : "primary",
            "unique" : true
          }
        ],
        "isSystem" : false,
        "isVolatile" : false,
        "journalSize" : 33554432,
        "keyOptions" : {
          "allowUserKeys" : true,
          "type" : "traditional"
        },
        "name" : "demo",
        "numberOfShards" : 1,
        "planId" : "93",
        "replicationFactor" : 1,
        "shardKeys" : [
          "_key"
        ],
        "shards" : {
        },
        "status" : 0,
        "type" : 2,
        "version" : 6,
        "waitForSync" : false
      },
      "indexes" : [ ]
    },
    {
      "parameters" : {
        "allowUserKeys" : true,
        "cid" : "10553",
        "count" : 0,
        "deleted" : false,
        "doCompact" : true,
        "id" : "10553",
        "indexBuckets" : 8,
        "indexes" : [
          {
            "fields" : [
              "_key"
            ],
            "id" : "0",
            "selectivityEstimate" : 1,
            "sparse" : false,
            "type" : "primary",
            "unique" : true
          },
          {
            "fields" : [
              "name"
            ],
            "id" : "10556",
            "selectivityEstimate" : 1,
            "sparse" : false,
            "type" : "hash",
            "unique" : false
          },
          {
            "fields" : [
              "a",
              "b"
            ],
            "id" : "10559",
            "sparse" : false,
            "type" : "skiplist",
            "unique" : true
          }
        ],
        "isSystem" : false,
        "isVolatile" : false,
        "journalSize" : 33554432,
        "keyOptions" : {
          "allowUserKeys" : true,
```

```
          "type" : "traditional"
        },
        "name" : "IndexedCollection1",
        "numberOfShards" : 1,
        "planId" : "10553",
        "replicationFactor" : 1,
        "shardKeys" : [
          "_key"
        ],
        "shards" : {
        },
        "status" : 3,
        "type" : 2,
        "version" : 6,
        "waitForSync" : false
      },
      "indexes" : [
        {
          "fields" : [
            "name"
          ],
          "id" : "10556",
          "selectivityEstimate" : 1,
          "sparse" : false,
          "type" : "hash",
          "unique" : false
        },
        {
          "fields" : [
            "a",
            "b"
          ],
          "id" : "10559",
          "sparse" : false,
          "type" : "skiplist",
          "unique" : true
        }
      ]
    },
    {
      "parameters" : {
        "allowUserKeys" : true,
        "cid" : "10561",
        "count" : 0,
        "deleted" : false,
        "doCompact" : true,
        "id" : "10561",
        "indexBuckets" : 8,
        "indexes" : [
          {
            "fields" : [
              "_key"
            ],
            "id" : "0",
            "selectivityEstimate" : 1,
            "sparse" : false,
            "type" : "primary",
            "unique" : true
          },
          {
            "fields" : [
              "text"
            ],
            "id" : "10564",
            "minLength" : 10,
            "sparse" : true,
            "type" : "fulltext",
            "unique" : false
          },
          {
            "fields" : [
              "a"
            ],
            "id" : "10567",
            "sparse" : false,
            "type" : "skiplist",
```

```
            "unique" : false
          }
        ],
        "isSystem" : false,
        "isVolatile" : false,
        "journalSize" : 33554432,
        "keyOptions" : {
          "allowUserKeys" : true,
          "type" : "traditional"
        },
        "name" : "IndexedCollection2",
        "numberOfShards" : 1,
        "planId" : "10561",
        "replicationFactor" : 1,
        "shardKeys" : [
          "_key"
        ],
        "shards" : {
        },
        "status" : 3,
        "type" : 2,
        "version" : 6,
        "waitForSync" : false
      },
      "indexes" : [
        {
          "fields" : [
            "text"
          ],
          "id" : "10564",
          "minLength" : 10,
          "sparse" : true,
          "type" : "fulltext",
          "unique" : false
        },
        {
          "fields" : [
            "a"
          ],
          "id" : "10567",
          "sparse" : false,
          "type" : "skiplist",
          "unique" : false
        }
      ]
    }
  ],
  "state" : {
    "running" : true,
    "lastLogTick" : "10565",
    "lastUncommittedLogTick" : "10568",
    "totalEvents" : 3811,
    "time" : "2017-04-18T15:53:17Z"
  },
  "tick" : "10568"
}
```

Return Codes

- *200*: is returned if the request was executed successfully.
- *405*: is returned when an invalid HTTP method is used.
- *500*: is returned if an error occurred while assembling the response.

Examples

```
shell> curl --dump - http://localhost:8529/_api/replication/inventory

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body
With some additional indexes:

```
shell> curl --dump - http://localhost:8529/_api/replication/inventory

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body
The *dump* method can be used to fetch data from a specific collection. As the results of the dump command can be huge, *dump* may not return all data from a collection at once. Instead, the dump command may be called repeatedly by replication clients until there is no more data to fetch. The dump command will not only return the current documents in the collection, but also document updates and deletions.

Please note that the *dump* method will only return documents, updates and deletions from a collection's journals and datafiles. Operations that are stored in the write-ahead log only will not be returned. In order to ensure that these operations are included in a dump, the write-ahead log must be flushed first.

To get to an identical state of data, replication clients should apply the individual parts of the dump results in the same order as they are provided.

Return data of a collection

returns the whole content of one collection

`GET /_api/replication/dump`

Query Parameters

- *collection* (required): The name or id of the collection to dump.
- *from* (optional): Lower bound tick value for results.
- *to* (optional): Upper bound tick value for results.
- *chunkSize* (optional): Approximate maximum size of the returned result.
- *includeSystem* (optional): Include system collections in the result. The default value is *true*.
- *failOnUnknown* (optional): Produce an error when dumped edges refer to now-unknown collections.
- *ticks* (optional): Whether or not to include tick values in the dump. The default value is *true*.
- *flush* (optional): Whether or not to flush the WAL before dumping. The default value is *true*.

Returns the data from the collection for the requested range.

When the *from* query parameter is not used, collection events are returned from the beginning. When the *from* parameter is used, the result will only contain collection entries which have higher tick values than the specified *from* value (note: the log entry with a tick value equal to *from* will be excluded).

The *to* query parameter can be used to optionally restrict the upper bound of the result to a certain tick value. If used, the result will only contain collection entries with tick values up to (including) *to*.

The *chunkSize* query parameter can be used to control the size of the result. It must be specified in bytes. The *chunkSize* value will only be honored approximately. Otherwise a too low *chunkSize* value could cause the server to not be able to put just one entry into the result and return it. Therefore, the *chunkSize* value will only be consulted after an entry has been written into the result. If the result size is then bigger than *chunkSize*, the server will respond with as many entries as there are in the response already. If the result size is still smaller than *chunkSize*, the server will try to return more data if there's more data left to return.

If *chunkSize* is not specified, some server-side default value will be used.

The *Content-Type* of the result is *application/x-arango-dump*. This is an easy-to-process format, with all entries going onto separate lines in the response body.

Each line itself is a JSON object, with at least the following attributes:

- *tick*: the operation's tick attribute

- *key*: the key of the document/edge or the key used in the deletion operation
- *rev*: the revision id of the document/edge or the deletion operation
- *data*: the actual document/edge data for types 2300 and 2301. The full document/edge data will be returned even for updates.
- *type*: the type of entry. Possible values for *type* are:
  - 2300: document insertion/update
  - 2301: edge insertion/update
  - 2302: document/edge deletion

Note: there will be no distinction between inserts and updates when calling this method.

Example:

Empty collection:

```
shell> curl --dump - http://localhost:8529/_api/replication/dump?collection=testCollection

HTTP/1.1 204 No Content
content-type: application/x-arango-dump; charset=utf-8
x-arango-replication-lastincluded: 0
x-arango-replication-checkmore: false
```

Example:

Non-empty collection:

```
shell> curl --dump - http://localhost:8529/_api/replication/dump?collection=testCollection

HTTP/1.1 200 OK
content-type: application/x-arango-dump; charset=utf-8
x-arango-replication-lastincluded: 10512
x-arango-replication-checkmore: false

"{\"tick\":\"10506\",\"type\":2300,\"data\":{\"_id\":\"testCollection/123456\",\"_key\":\"123456\",\"_rev\":\"_U2FCPw
y--_\",\"b\":1,\"c\":false,\"d\":\"additional value\"}}
{\"tick\":\"10510\",\"type\":2302,\"data\":{\"_key\":\"foobar\",\"_rev\":\"_U2FCPw6---\"}}
{\"tick\":\"10512\",\"type\":2302,\"data\":{\"_key\":\"abcdef\",\"_rev\":\"_U2FCPw6--_\"}}
"
```

Return Codes

- *200*: is returned if the request was executed successfully and data was returned. The header `x-arango-replication-lastincluded` is set to the tick of the last document returned.
- *204*: is returned if the request was executed successfully, but there was no content available. The header `x-arango-replication-lastincluded` is `0` in this case.
- *400*: is returned if either the *from* or *to* values are invalid.
- *404*: is returned when the collection could not be found.
- *405*: is returned when an invalid HTTP method is used.
- *500*: is returned if an error occurred while assembling the response.

Examples

Empty collection:

```
shell> curl --dump - http://localhost:8529/_api/replication/dump?
collection=testCollection

HTTP/1.1 204 No Content
content-type: application/x-arango-dump; charset=utf-8
x-arango-replication-lastincluded: 0
x-arango-replication-checkmore: false
```

Non-empty collection:

```
shell> curl --dump - http://localhost:8529/_api/replication/dump?
collection=testCollection

HTTP/1.1 200 OK
content-type: application/x-arango-dump; charset=utf-8
x-arango-replication-lastincluded: 10512
x-arango-replication-checkmore: false

"{\"tick\":\"10506\",\"type\":2300,\"data\":
{\"_id\":\"testCollection/123456\",\"_key\":\"123456\",\"_rev\":\"_U2FCPwy--
_\",\"b\":1,\"c\":false,\"d\":\"additional
value\"}}\n{\"tick\":\"10510\",\"type\":2302,\"data\":
{\"_key\":\"foobar\",\"_rev\":\"_U2FCPw6---
\"}}\n{\"tick\":\"10512\",\"type\":2302,\"data\":
{\"_key\":\"abcdef\",\"_rev\":\"_U2FCPw6--_\"}}\n"
```

Synchronize data from a remote endpoint

start a replication

`PUT /_api/replication/sync`

A json post document with these Properties is required:

- **username**: an optional ArangoDB username to use when connecting to the endpoint.
- **includeSystem**: whether or not system collection operations will be applied
- **endpoint**: the master endpoint to connect to (e.g. "tcp://192.168.173.13:8529").
- **initialSyncMaxWaitTime**: the maximum wait time (in seconds) that the initial synchronization will wait for a response from the master when fetching initial collection data. This wait time can be used to control after what time the initial synchronization will give up waiting for a response and fail. This value will be ignored if set to *0*.
- **database**: the database name on the master (if not specified, defaults to the name of the local current database).
- **restrictType**: an optional string value for collection filtering. When specified, the allowed values are *include* or *exclude*.
- **incremental**: if set to *true*, then an incremental synchronization method will be used for synchronizing data in collections. This method is useful when collections already exist locally, and only the remaining differences need to be transferred from the remote endpoint. In this case, the incremental synchronization can be faster than a full synchronization. The default value is *false*, meaning that the complete data from the remote collection will be transferred.
- **restrictCollections** (string): an optional array of collections for use with *restrictType*. If *restrictType* is *include*, only the specified collections will be sychronised. If *restrictType* is *exclude*, all but the specified collections will be synchronized.
- **password**: the password to use when connecting to the endpoint.

Starts a full data synchronization from a remote endpoint into the local ArangoDB database.

The *sync* method can be used by replication clients to connect an ArangoDB database to a remote endpoint, fetch the remote list of collections and indexes, and collection data. It will thus create a local backup of the state of data at the remote ArangoDB database. *sync* works on a per-database level.

sync will first fetch the list of collections and indexes from the remote endpoint. It does so by calling the *inventory* API of the remote database. It will then purge data in the local ArangoDB database, and after start will transfer collection data from the remote database to the local ArangoDB database. It will extract data from the remote database by calling the remote database's *dump* API until all data are fetched.

In case of success, the body of the response is a JSON object with the following attributes:

- *collections*: an array of collections that were transferred from the endpoint
- *lastLogTick*: the last log tick on the endpoint at the time the transfer was started. Use this value as the *from* value when starting the continuous synchronization later.

WARNING: calling this method will sychronize data from the collections found on the remote endpoint to the local ArangoDB database. All data in the local collections will be purged and replaced with data from the endpoint.

Use with caution!

Note: this method is not supported on a coordinator in a cluster.

Return Codes

- *200*: is returned if the request was executed successfully.
- *400*: is returned if the configuration is incomplete or malformed.
- *405*: is returned when an invalid HTTP method is used.
- *500*: is returned if an error occurred during sychronization.
- *501*: is returned when this operation is called on a coordinator in a cluster.

Return cluster inventory of collections and indexes

returs an overview of collections and indexes in a cluster

`GET /_api/replication/clusterInventory`

Query Parameters

- *includeSystem* (optional): Include system collections in the result. The default value is *true*.

Returns the array of collections and indexes available on the cluster.

The response will be an array of JSON objects, one for each collection. Each collection containscontains exactly two keys "parameters" and "indexes". This information comes from Plan/Collections/{DB-Name}/ *in the agency, just that the* indexes* attribute there is relocated to adjust it to the data format of arangodump.

Return Codes

- *200*: is returned if the request was executed successfully.
- *405*: is returned when an invalid HTTP method is used.
- *500*: is returned if an error occurred while assembling the response.

Replication Logger Commands

Previous versions of ArangoDB allowed starting, stopping and configuring the replication logger. These commands are superfluous in ArangoDB 2.2 as all data-modification operations are written to the server's write-ahead log and are not handled by a separate logger anymore.

The only useful operations remaining since ArangoDB 2.2 are to query the current state of the logger and to fetch the latest changes written by the logger. The operations will return the state and data from the write-ahead log.

Return replication logger state

returns the state of the replication logger

`GET /_api/replication/logger-state`

Returns the current state of the server's replication logger. The state will include information about whether the logger is running and about the last logged tick value. This tick value is important for incremental fetching of data.

The body of the response contains a JSON object with the following attributes:

- *state*: the current logger state as a JSON object with the following sub-attributes:
  - *running*: whether or not the logger is running
  - *lastLogTick*: the tick value of the latest tick the logger has logged. This value can be used for incremental fetching of log data.
  - *totalEvents*: total number of events logged since the server was started. The value is not reset between multiple stops and re-starts of the logger.
  - *time*: the current date and time on the logger server
- *server*: a JSON object with the following sub-attributes:
  - *version*: the logger server's version
  - *serverId*: the logger server's id
- *clients*: returns the last fetch status by replication clients connected to the logger. Each client is returned as a JSON object with the following attributes:
  - *serverId*: server id of client
  - *lastServedTick*: last tick value served to this client via the *logger-follow* API
  - *time*: date and time when this client last called the *logger-follow* API

Example:

Returns the state of the replication logger.

```
shell> curl --dump - http://localhost:8529/_api/replication/logger-state

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "state" : {
    "running" : true,
    "lastLogTick" : "10594",
    "lastUncommittedLogTick" : "10594",
    "totalEvents" : 3825,
    "time" : "2017-04-18T15:53:20Z"
  },
  "server" : {
    "version" : "3.1.18",
    "serverId" : "271514875687090"
  },
  "clients" : [ ]
}
```

Return Codes

- *200*: is returned if the logger state could be determined successfully.
- *405*: is returned when an invalid HTTP method is used.
- *500*: is returned if the logger state could not be determined.

Examples

Returns the state of the replication logger.

```
shell> curl --dump - http://localhost:8529/_api/replication/logger-state

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

To query the latest changes logged by the replication logger, the HTTP interface also provides the `logger-follow` method.

This method should be used by replication clients to incrementally fetch updates from an ArangoDB database.

Returns log entries

Fetch log lines from the server

`GET /_api/replication/logger-follow`

Query Parameters

- *from* (optional): Lower bound tick value for results.
- *to* (optional): Upper bound tick value for results.
- *chunkSize* (optional): Approximate maximum size of the returned result.
- *includeSystem* (optional): Include system collections in the result. The default value is *true*.

Returns data from the server's replication log. This method can be called by replication clients after an initial synchronization of data. The method will return all "recent" log entries from the logger server, and the clients can replay and apply these entries locally so they get to the same data state as the logger server.

Clients can call this method repeatedly to incrementally fetch all changes from the logger server. In this case, they should provide the *from* value so they will only get returned the log events since their last fetch.

When the *from* query parameter is not used, the logger server will return log entries starting at the beginning of its replication log. When the *from* parameter is used, the logger server will only return log entries which have higher tick values than the specified *from* value (note: the log entry with a tick value equal to *from* will be excluded). Use the *from* value when incrementally fetching log data.

The *to* query parameter can be used to optionally restrict the upper bound of the result to a certain tick value. If used, the result will contain only log events with tick values up to (including) *to*. In incremental fetching, there is no need to use the *to* parameter. It only makes sense in special situations, when only parts of the change log are required.

The *chunkSize* query parameter can be used to control the size of the result. It must be specified in bytes. The *chunkSize* value will only be honored approximately. Otherwise a too low *chunkSize* value could cause the server to not be able to put just one log entry into the result and return it. Therefore, the *chunkSize* value will only be consulted after a log entry has been written into the result. If the result size is then bigger than *chunkSize*, the server will respond with as many log entries as there are in the response already. If the result size is still smaller than *chunkSize*, the server will try to return more data if there's more data left to return.

If *chunkSize* is not specified, some server-side default value will be used.

The *Content-Type* of the result is *application/x-arango-dump*. This is an easy-to-process format, with all log events going onto separate lines in the response body. Each log event itself is a JSON object, with at least the following attributes:

- *tick*: the log event tick value
- *type*: the log event type

Individual log events will also have additional attributes, depending on the event type. A few common attributes which are used for multiple events types are:

- *cid*: id of the collection the event was for
- *tid*: id of the transaction the event was contained in
- *key*: document key
- *rev*: document revision id
- *data*: the original document data

A more detailed description of the individual replication event types and their data structures can be found in the manual.

The response will also contain the following HTTP headers:

- *x-arango-replication-active*: whether or not the logger is active. Clients can use this flag as an indication for their polling frequency. If the logger is not active and there are no more replication events available, it might be sensible for a client to abort, or to go to sleep for a long time and try again later to check whether the logger has been activated.
- *x-arango-replication-lastincluded*: the tick value of the last included value in the result. In incremental log fetching, this value can be used as the *from* value for the following request. **Note** that if the result is empty, the value will be *0*. This value should not be used as *from* value by clients in the next request (otherwise the server would return the log events from the start of the log again).
- *x-arango-replication-lasttick*: the last tick value the logger server has logged (not necessarily included in the result). By comparing the the last tick and last included tick values, clients have an approximate indication of how many events there are still left to fetch.
- *x-arango-replication-checkmore*: whether or not there already exists more log data which the client could fetch immediately. If there is more log data available, the client could call *logger-follow* again with an adjusted *from* value to fetch remaining log entries until there are no more.

  If there isn't any more log data to fetch, the client might decide to go to sleep for a while before calling the logger again.

Note: this method is not supported on a coordinator in a cluster.

Example:

No log events available

```
shell> curl --dump - http://localhost:8529/_api/replication/logger-follow?from=10582

HTTP/1.1 204 No Content
content-type: application/x-arango-dump; charset=utf-8
x-arango-replication-lasttick: 10582
x-arango-replication-active: true
x-arango-replication-frompresent: true
x-arango-replication-checkmore: false
x-arango-replication-lastincluded: 0
```

Example:

A few log events

```
shell> curl --dump - http://localhost:8529/_api/replication/logger-follow?from=10582

HTTP/1.1 200 OK
content-type: application/x-arango-dump; charset=utf-8
x-arango-replication-lasttick: 10594
x-arango-replication-active: true
x-arango-replication-frompresent: true
x-arango-replication-checkmore: false
x-arango-replication-lastincluded: 10594

"{\"tick\":\"10584\",\"type\":2000,\"database\":\"1\",\"cid\":\"10583\",\"cname\":\"products\",\"data\":{\"allowUserK
eys\":true,\"cid\":\"10583\",\"count\":0,\"deleted\":false,\"doCompact\":true,\"id\":\"10583\",\"indexBuckets\":8,\"i
ndexes\":[{\"fields\":[\"_key\"],\"id\":\"0\",\"selectivityEstimate\":1,\"sparse\":false,\"type\":\"primary\",\"uniqu
e\":true}],\"isSystem\":false,\"isVolatile\":false,\"journalSize\":33554432,\"keyOptions\":{\"allowUserKeys\":true,\"
type\":\"traditional\"},\"name\":\"products\",\"numberOfShards\":1,\"planId\":\"10583\",\"replicationFactor\":1,\"sha
rdKeys\":[\"_key\"],\"shards\":{},\"status\":3,\"type\":2,\"version\":6,\"waitForSync\":false}}
{\"tick\":\"10587\",\"type\":2300,\"tid\":\"0\",\"database\":\"1\",\"cid\":\"10583\",\"cname\":\"products\",\"data\":
{\"_id\":\"_unknown/p1\",\"_key\":\"p1\",\"_rev\":\"_U2FCmVC---\",\"name\":\"flux compensator\"}}
{\"tick\":\"10589\",\"type\":2300,\"tid\":\"0\",\"database\":\"1\",\"cid\":\"10583\",\"cname\":\"products\",\"data\":
{\"_id\":\"_unknown/p2\",\"_key\":\"p2\",\"_rev\":\"_U2FCmVC--_\",\"hp\":5100,\"name\":\"hybrid hovercraft\"}}
{\"tick\":\"10591\",\"type\":2302,\"tid\":\"0\",\"database\":\"1\",\"cid\":\"10583\",\"cname\":\"products\",\"data\":
{\"_key\":\"p1\",\"_rev\":\"_U2FCmVG---\"}}
{\"tick\":\"10593\",\"type\":2300,\"tid\":\"0\",\"database\":\"1\",\"cid\":\"10583\",\"cname\":\"products\",\"data\":
{\"_id\":\"_unknown/p2\",\"_key\":\"p2\",\"_rev\":\"_U2FCmVG--_\",\"hp\":5100,\"name\":\"broken hovercraft\"}}
{\"tick\":\"10594\",\"type\":2001,\"database\":\"1\",\"cid\":\"10583\",\"cname\":\"products\",\"data\":{\"id\":\"1058
3\",\"name\":\"products\"}}
"
```

Example:

More events than would fit into the response

```
shell> curl --dump - http://localhost:8529/_api/replication/logger-follow?from=10570&chunkSize=400

HTTP/1.1 200 OK
content-type: application/x-arango-dump; charset=utf-8
x-arango-replication-lasttick: 10582
x-arango-replication-active: true
x-arango-replication-frompresent: true
x-arango-replication-checkmore: true
x-arango-replication-lastincluded: 10572

"{\"tick\":\"10572\",\"type\":2000,\"database\":\"1\",\"cid\":\"10571\",\"cname\":\"products\",\"data\":{\"allowUserK
eys\":true,\"cid\":\"10571\",\"count\":0,\"deleted\":false,\"doCompact\":true,\"id\":\"10571\",\"indexBuckets\":8,\"i
ndexes\":[{\"fields\":[\"_key\"],\"id\":\"0\",\"selectivityEstimate\":1,\"sparse\":false,\"type\":\"primary\",\"uniqu
e\":true}],\"isSystem\":false,\"isVolatile\":false,\"journalSize\":33554432,\"keyOptions\":{\"allowUserKeys\":true,\"
type\":\"traditional\"},\"name\":\"products\",\"numberOfShards\":1,\"planId\":\"10571\",\"replicationFactor\":1,\"sha
rdKeys\":[\"_key\"],\"shards\":{},\"status\":3,\"type\":2,\"version\":6,\"waitForSync\":false}}
"
```

Return Codes

- *200*: is returned if the request was executed successfully, and there are log events available for the requested range. The response body will not be empty in this case.
- *204*: is returned if the request was executed successfully, but there are no log events available for the requested range. The response body will be empty in this case.
- *400*: is returned if either the *from* or *to* values are invalid.
- *405*: is returned when an invalid HTTP method is used.
- *500*: is returned if an error occurred while assembling the response.
- *501*: is returned when this operation is called on a coordinator in a cluster.

Examples

No log events available

```
shell> curl --dump - http://localhost:8529/_api/replication/logger-follow?from=10582

HTTP/1.1 204 No Content
content-type: application/x-arango-dump; charset=utf-8
x-arango-replication-lasttick: 10582
x-arango-replication-active: true
x-arango-replication-frompresent: true
x-arango-replication-checkmore: false
x-arango-replication-lastincluded: 0
```

A few log events

```
shell> curl --dump - http://localhost:8529/_api/replication/logger-follow?from=10582

HTTP/1.1 200 OK
content-type: application/x-arango-dump; charset=utf-8
x-arango-replication-lasttick: 10594
x-arango-replication-active: true
x-arango-replication-frompresent: true
x-arango-replication-checkmore: false
x-arango-replication-lastincluded: 10594

"
{\"tick\":\"10584\",\"type\":2000,\"database\":\"1\",\"cid\":\"10583\",\"cname\":\"produ
cts\",\"data\":
{\"allowUserKeys\":true,\"cid\":\"10583\",\"count\":0,\"deleted\":false,\"doCompact\":tr
ue,\"id\":\"10583\",\"indexBuckets\":8,\"indexes\":[{\"fields\":
[\"_key\"],\"id\":\"0\",\"selectivityEstimate\":1,\"sparse\":false,\"type\":\"primary\",
\"unique\":true}],\"isSystem\":false,\"isVolatile\":false,\"journalSize\":33554432,\"key
Options\":
{\"allowUserKeys\":true,\"type\":\"traditional\"},\"name\":\"products\",\"numberOfShards
\":1,\"planId\":\"10583\",\"replicationFactor\":1,\"shardKeys\":[\"_key\"],\"shards\":
{},\"status\":3,\"type\":2,\"version\":6,\"waitForSync\":false}}\n{\"tick\":\"10587\",\"
type\":2300,\"tid\":\"0\",\"database\":\"1\",\"cid\":\"10583\",\"cname\":\"products\",\"
data\":{\"_id\":\"_unknown/p1\",\"_key\":\"p1\",\"_rev\":\"_U2FCmVC---\",\"name\":\"flux
compensator\"}}\n{\"tick\":\"10589\",\"type\":2300,\"tid\":\"0\",\"database\":\"1\",\"ci
d\":\"10583\",\"cname\":\"products\",\"data\":
{\"_id\":\"_unknown/p2\",\"_key\":\"p2\",\"_rev\":\"_U2FCmVC--
_\",\"hp\":5100,\"name\":\"hybrid
hovercraft\"}}\n{\"tick\":\"10591\",\"type\":2302,\"tid\":\"0\",\"database\":\"1\",\"cid
\":\"10583\",\"cname\":\"products\",\"data\":{\"_key\":\"p1\",\"_rev\":\"_U2FCmVG---
\"}}\n{\"tick\":\"10593\",\"type\":2300,\"tid\":\"0\",\"database\":\"1\",\"cid\":\"10583
\",\"cname\":\"products\",\"data\":
{\"_id\":\"_unknown/p2\",\"_key\":\"p2\",\"_rev\":\"_U2FCmVG--
_\",\"hp\":5100,\"name\":\"broken
hovercraft\"}}\n{\"tick\":\"10594\",\"type\":2001,\"database\":\"1\",\"cid\":\"10583\",\
"cname\":\"products\",\"data\":{\"id\":\"10583\",\"name\":\"products\"}}\n"
```

More events than would fit into the response

```
shell> curl --dump - http://localhost:8529/_api/replication/logger-follow?
from=10570&chunkSize=400

HTTP/1.1 200 OK
content-type: application/x-arango-dump; charset=utf-8
x-arango-replication-lasttick: 10582
x-arango-replication-active: true
x-arango-replication-frompresent: true
x-arango-replication-checkmore: true
x-arango-replication-lastincluded: 10572

"
{\"tick\":\"10572\",\"type\":2000,\"database\":\"1\",\"cid\":\"10571\",\"cname\":\"produ
cts\",\"data\":
{\"allowUserKeys\":true,\"cid\":\"10571\",\"count\":0,\"deleted\":false,\"doCompact\":tr
ue,\"id\":\"10571\",\"indexBuckets\":8,\"indexes\":[{\"fields\":
[\"_key\"],\"id\":\"0\",\"selectivityEstimate\":1,\"sparse\":false,\"type\":\"primary\",
\"unique\":true}],\"isSystem\":false,\"isVolatile\":false,\"journalSize\":33554432,\"key
Options\":
{\"allowUserKeys\":true,\"type\":\"traditional\"},\"name\":\"products\",\"numberOfShards
\":1,\"planId\":\"10571\",\"replicationFactor\":1,\"shardKeys\":[\"_key\"],\"shards\":
{},\"status\":3,\"type\":2,\"version\":6,\"waitForSync\":false}}\n"
```

To check what range of changes is available (identified by tick values), the HTTP interface provides the methods `logger-first-tick` and `logger-tick-ranges`. Replication clients can use the methods to determine if certain data (identified by a tick *date*) is still available on the master.

Returns the first available tick value

Return the first available tick value from the server

`GET /_api/replication/logger-first-tick`

Returns the first available tick value that can be served from the server's replication log. This method can be called by replication clients after to determine if certain data (identified by a tick value) is still available for replication.

The result is a JSON object containing the attribute *firstTick*. This attribute contains the minimum tick value available in the server's replication log.

Note: this method is not supported on a coordinator in a cluster.

Example:

Returning the first available tick

```
shell> curl --dump - http://localhost:8529/_api/replication/logger-first-tick

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

"{\"firstTick\":\"5\"}"
```

Return Codes

- *200*: is returned if the request was executed successfully.
- *405*: is returned when an invalid HTTP method is used.
- *500*: is returned if an error occurred while assembling the response.
- *501*: is returned when this operation is called on a coordinator in a cluster.

Examples

Returning the first available tick

```
shell> curl --dump - http://localhost:8529/_api/replication/logger-first-tick

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

"{\"firstTick\":\"5\"}"
```

Return the tick ranges available in the WAL logfiles

returns the tick value ranges available in the logfiles

`GET /_api/replication/logger-tick-ranges`

Returns the currently available ranges of tick values for all currently available WAL logfiles. The tick values can be used to determine if certain data (identified by tick value) are still available for replication.

The body of the response contains a JSON array. Each array member is an object that describes a single logfile. Each object has the following attributes:

- *datafile*: name of the logfile
- *status*: status of the datafile, in textual form (e.g. "sealed", "open")
- *tickMin*: minimum tick value contained in logfile
- *tickMax*: maximum tick value contained in logfile

Example:

Returns the available tick ranges.

```
shell> curl --dump - http://localhost:8529/_api/replication/logger-tick-ranges

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

[
  {
    "datafile" : "/tmp/arangosh_GgAPmL/tmp-15793-280372105/data/journals/logfile-2.db",
    "status" : "collected",
    "tickMin" : "5",
    "tickMax" : "9453"
  },
  {
    "datafile" : "/tmp/arangosh_GgAPmL/tmp-15793-280372105/data/journals/logfile-23.db",
    "status" : "collected",
    "tickMin" : "9471",
    "tickMax" : "9512"
  },
  {
    "datafile" : "/tmp/arangosh_GgAPmL/tmp-15793-280372105/data/journals/logfile-49.db",
    "status" : "collected",
    "tickMin" : "9518",
    "tickMax" : "10512"
  },
  {
    "datafile" : "/tmp/arangosh_GgAPmL/tmp-15793-280372105/data/journals/logfile-9456.db",
    "status" : "collected",
    "tickMin" : "10519",
    "tickMax" : "10545"
  },
  {
    "datafile" : "/tmp/arangosh_GgAPmL/tmp-15793-280372105/data/journals/logfile-9515.db",
    "status" : "open",
    "tickMin" : "10549",
    "tickMax" : "10594"
  }
]
```

Return Codes

- *200*: is returned if the tick ranges could be determined successfully.
- *405*: is returned when an invalid HTTP method is used.
- *500*: is returned if the logger state could not be determined.
- *501*: is returned when this operation is called on a coordinator in a cluster.

Examples

Returns the available tick ranges.

```
shell> curl --dump - http://localhost:8529/_api/replication/logger-tick-ranges

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

[
  {
    "datafile" : "/tmp/arangosh_GgAPmL/tmp-15793-280372105/data/journals/logfile-2.db",
    "status" : "collected",
    "tickMin" : "5",
    "tickMax" : "9453"
  },
  {
    "datafile" : "/tmp/arangosh_GgAPmL/tmp-15793-280372105/data/journals/logfile-23.db",
    "status" : "collected",
    "tickMin" : "9471",
    "tickMax" : "9512"
  },
  {
    "datafile" : "/tmp/arangosh_GgAPmL/tmp-15793-280372105/data/journals/logfile-49.db",
    "status" : "collected",
    "tickMin" : "9518",
    "tickMax" : "10512"
  },
  {
    "datafile" : "/tmp/arangosh_GgAPmL/tmp-15793-280372105/data/journals/logfile-
9456.db",
    "status" : "collected",
    "tickMin" : "10519",
    "tickMax" : "10545"
  },
  {
    "datafile" : "/tmp/arangosh_GgAPmL/tmp-15793-280372105/data/journals/logfile-
9515.db",
    "status" : "open",
    "tickMin" : "10549",
    "tickMax" : "10594"
  }
]
```

Replication Applier Commands

The applier commands allow to remotely start, stop, and query the state and configuration of an ArangoDB database's replication applier.

Return configuration of replication applier

fetch the current replication configuration

`GET /_api/replication/applier-config`

Returns the configuration of the replication applier.

The body of the response is a JSON object with the configuration. The following attributes may be present in the configuration:

- *endpoint*: the logger server to connect to (e.g. "tcp://192.168.173.13:8529").
- *database*: the name of the database to connect to (e.g. "_system").
- *username*: an optional ArangoDB username to use when connecting to the endpoint.
- *password*: the password to use when connecting to the endpoint.
- *maxConnectRetries*: the maximum number of connection attempts the applier will make in a row. If the applier cannot establish a connection to the endpoint in this number of attempts, it will stop itself.
- *connectTimeout*: the timeout (in seconds) when attempting to connect to the endpoint. This value is used for each connection attempt.
- *requestTimeout*: the timeout (in seconds) for individual requests to the endpoint.
- *chunkSize*: the requested maximum size for log transfer packets that is used when the endpoint is contacted.
- *autoStart*: whether or not to auto-start the replication applier on (next and following) server starts
- *adaptivePolling*: whether or not the replication applier will use adaptive polling.
- *includeSystem*: whether or not system collection operations will be applied
- *autoResync*: whether or not the slave should perform a full automatic resynchronization with the master in case the master cannot serve log data requested by the slave, or when the replication is started and no tick value can be found.
- *autoResyncRetries*: number of resynchronization retries that will be performed in a row when automatic resynchronization is enabled and kicks in. Setting this to *0* will effectively disable *autoResync*. Setting it to some other value will limit the number of retries that are performed. This helps preventing endless retries in case resynchronizations always fail.
- *initialSyncMaxWaitTime*: the maximum wait time (in seconds) that the initial synchronization will wait for a response from the master when fetching initial collection data. This wait time can be used to control after what time the initial synchronization will give up waiting for a response and fail. This value is relevant even for continuous replication when *autoResync* is set to *true* because this may re-start the initial synchronization when the master cannot provide log data the slave requires. This value will be ignored if set to *0*.
- *connectionRetryWaitTime*: the time (in seconds) that the applier will intentionally idle before it retries connecting to the master in case of connection problems. This value will be ignored if set to *0*.
- *idleMinWaitTime*: the minimum wait time (in seconds) that the applier will intentionally idle before fetching more log data from the master in case the master has already sent all its log data. This wait time can be used to control the frequency with which the replication applier sends HTTP log fetch requests to the master in case there is no write activity on the master. This value will be ignored if set to *0*.
- *idleMaxWaitTime*: the maximum wait time (in seconds) that the applier will intentionally idle before fetching more log data from the master in case the master has already sent all its log data and there have been previous log fetch attempts that resulted in no more log data. This wait time can be used to control the maximum frequency with which the replication applier sends HTTP log fetch requests to the master in case there is no write activity on the master for longer periods. This configuration value will only be used if the option *adaptivePolling* is set to *true*. This value will be ignored if set to *0*.

- *requireFromPresent*: if set to *true*, then the replication applier will check at start whether the start tick from which it starts or resumes replication is still present on the master. If not, then there would be data loss. If *requireFromPresent* is *true*, the replication applier will abort with an appropriate error message. If set to *false*, then the replication applier will still start, and ignore the data loss.
- *verbose*: if set to *true*, then a log line will be emitted for all operations performed by the replication applier. This should be used for debugging replication problems only.
- *restrictType*: the configuration for *restrictCollections*
- *restrictCollections*: the optional array of collections to include or exclude, based on the setting of *restrictType*

Example:

```
shell> curl --dump - http://localhost:8529/_api/replication/applier-config

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "requestTimeout" : 600,
  "connectTimeout" : 10,
  "ignoreErrors" : 0,
  "maxConnectRetries" : 100,
  "sslProtocol" : 0,
  "chunkSize" : 0,
  "autoStart" : false,
  "adaptivePolling" : true,
  "autoResync" : false,
  "autoResyncRetries" : 2,
  "includeSystem" : true,
  "requireFromPresent" : false,
  "verbose" : false,
  "incremental" : false,
  "useCollectionId" : true,
  "restrictType" : "",
  "restrictCollections" : [ ],
  "connectionRetryWaitTime" : 15,
  "initialSyncMaxWaitTime" : 300,
  "idleMinWaitTime" : 1,
  "idleMaxWaitTime" : 2.5
}
```

Return Codes

- *200*: is returned if the request was executed successfully.
- *405*: is returned when an invalid HTTP method is used.
- *500*: is returned if an error occurred while assembling the response.

Examples

```
shell> curl --dump - http://localhost:8529/_api/replication/applier-config

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Adjust configuration of replication applier

set configuration values of an applier

`PUT /_api/replication/applier-config`

A json post document with these Properties is required:

- **username**: an optional ArangoDB username to use when connecting to the endpoint.
- **includeSystem**: whether or not system collection operations will be applied
- **endpoint**: the logger server to connect to (e.g. "tcp://192.168.173.13:8529"). The endpoint must be specified.
- **verbose**: if set to *true*, then a log line will be emitted for all operations performed by the replication applier. This should be used for debugging replication problems only.
- **connectTimeout**: the timeout (in seconds) when attempting to connect to the endpoint. This value is used for each connection attempt.
- **autoResync**: whether or not the slave should perform a full automatic resynchronization with the master in case the master cannot serve log data requested by the slave, or when the replication is started and no tick value can be found.
- **database**: the name of the database on the endpoint. If not specified, defaults to the current local database name.
- **idleMinWaitTime**: the minimum wait time (in seconds) that the applier will intentionally idle before fetching more log data from the master in case the master has already sent all its log data. This wait time can be used to control the frequency with which the replication applier sends HTTP log fetch requests to the master in case there is no write activity on the master. This value will be ignored if set to *0*.
- **requestTimeout**: the timeout (in seconds) for individual requests to the endpoint.
- **requireFromPresent**: if set to *true*, then the replication applier will check at start whether the start tick from which it starts or resumes replication is still present on the master. If not, then there would be data loss. If *requireFromPresent* is *true*, the replication applier will abort with an appropriate error message. If set to *false*, then the replication applier will still start, and ignore the data loss.
- **idleMaxWaitTime**: the maximum wait time (in seconds) that the applier will intentionally idle before fetching more log data from the master in case the master has already sent all its log data and there have been previous log fetch attempts that resulted in no more log data. This wait time can be used to control the maximum frequency with which the replication applier sends HTTP log fetch requests to the master in case there is no write activity on the master for longer periods. This configuration value will only be used if the option *adaptivePolling* is set to *true*. This value will be ignored if set to *0*.
- **restrictCollections** (string): the array of collections to include or exclude, based on the setting of *restrictType*
- **restrictType**: the configuration for *restrictCollections*; Has to be either *include* or *exclude*
- **initialSyncMaxWaitTime**: the maximum wait time (in seconds) that the initial synchronization will wait for a response from the master when fetching initial collection data. This wait time can be used to control after what time the initial synchronization will give up waiting for a response and fail. This value is relevant even for continuous replication when *autoResync* is set to *true* because this may re-start the initial synchronization when the master cannot provide log data the slave requires. This value will be ignored if set to *0*.
- **maxConnectRetries**: the maximum number of connection attempts the applier will make in a row. If the applier cannot establish a connection to the endpoint in this number of attempts, it will stop itself.
- **autoStart**: whether or not to auto-start the replication applier on (next and following) server starts
- **adaptivePolling**: if set to *true*, the replication applier will fall to sleep for an increasingly long period in case the logger server at the endpoint does not have any more replication events to apply. Using adaptive polling is thus useful to reduce the amount of work for both the applier and the logger server for cases when there are only infrequent changes. The downside is that when using adaptive polling, it might take longer for the replication applier to detect that there are new replication events on the logger server. Setting *adaptivePolling* to false will make the replication applier contact the logger server in a constant interval, regardless of whether the logger server provides updates frequently or seldom.
- **password**: the password to use when connecting to the endpoint.
- **connectionRetryWaitTime**: the time (in seconds) that the applier will intentionally idle before it retries connecting to the master in case of connection problems. This value will be ignored if set to *0*.
- **autoResyncRetries**: number of resynchronization retries that will be performed in a row when automatic resynchronization is enabled and kicks in. Setting this to *0* will effectively disable *autoResync*. Setting it to some other value will limit the number of retries that are performed. This helps preventing endless retries in case resynchronizations always fail.
- **chunkSize**: the requested maximum size for log transfer packets that is used when the endpoint is contacted.

Sets the configuration of the replication applier. The configuration can only be changed while the applier is not running. The updated configuration will be saved immediately but only become active with the next start of the applier.

In case of success, the body of the response is a JSON object with the updated configuration.

Example:

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/replication/applier-config <<EOF
{
  "endpoint" : "tcp://127.0.0.1:8529",
  "username" : "replicationApplier",
  "password" : "applier1234@foxx",
  "chunkSize" : 4194304,
  "autoStart" : false,
  "adaptivePolling" : true
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "endpoint" : "tcp://127.0.0.1:8529",
  "database" : "_system",
  "username" : "replicationApplier",
  "requestTimeout" : 600,
  "connectTimeout" : 10,
  "ignoreErrors" : 0,
  "maxConnectRetries" : 100,
  "sslProtocol" : 0,
  "chunkSize" : 4194304,
  "autoStart" : false,
  "adaptivePolling" : true,
  "autoResync" : false,
  "autoResyncRetries" : 2,
  "includeSystem" : true,
  "requireFromPresent" : false,
  "verbose" : false,
  "incremental" : false,
  "useCollectionId" : true,
  "restrictType" : "",
  "restrictCollections" : [ ],
  "connectionRetryWaitTime" : 15,
  "initialSyncMaxWaitTime" : 300,
  "idleMinWaitTime" : 1,
  "idleMaxWaitTime" : 2.5
}
```

Return Codes

- *200*: is returned if the request was executed successfully.
- *400*: is returned if the configuration is incomplete or malformed, or if the replication applier is currently running.
- *405*: is returned when an invalid HTTP method is used.
- *500*: is returned if an error occurred while assembling the response.

Examples

```
shell> curl -X PUT --data-binary @- --dump -
http://localhost:8529/_api/replication/applier-config <<EOF
{
  "endpoint" : "tcp://127.0.0.1:8529",
  "username" : "replicationApplier",
  "password" : "applier1234@foxx",
  "chunkSize" : 4194304,
  "autoStart" : false,
  "adaptivePolling" : true
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Start replication applier

start the replication applier

`PUT /_api/replication/applier-start`

Query Parameters

- *from* (optional): The remote *lastLogTick* value from which to start applying. If not specified, the last saved tick from the previous applier run is used. If there is no previous applier state saved, the applier will start at the beginning of the logger server's log.

Starts the replication applier. This will return immediately if the replication applier is already running.

If the replication applier is not already running, the applier configuration will be checked, and if it is complete, the applier will be started in a background thread. This means that even if the applier will encounter any errors while running, they will not be reported in the response to this method.

To detect replication applier errors after the applier was started, use the */_api/replication/applier-state* API instead.

Example:

```
shell> curl -X PUT --dump - http://localhost:8529/_api/replication/applier-start

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "state" : {
    "running" : true,
    "lastAppliedContinuousTick" : null,
    "lastProcessedContinuousTick" : null,
    "lastAvailableContinuousTick" : null,
    "safeResumeTick" : null,
    "progress" : {
      "time" : "2017-04-18T15:51:40Z",
      "message" : "applier initially created",
      "failedConnects" : 0
    },
    "totalRequests" : 0,
    "totalFailedConnects" : 0,
    "totalEvents" : 0,
    "totalOperationsExcluded" : 0,
    "lastError" : {
      "errorNum" : 0
    },
    "time" : "2017-04-18T15:52:54Z"
  },
  "server" : {
    "version" : "3.1.18",
    "serverId" : "271514875687090"
  },
  "endpoint" : "tcp://127.0.0.1:8529",
  "database" : "_system"
}
```

Return Codes

- *200*: is returned if the request was executed successfully.
- *400*: is returned if the replication applier is not fully configured or the configuration is invalid.
- *405*: is returned when an invalid HTTP method is used.
- *500*: is returned if an error occurred while assembling the response.

Examples

```
shell> curl -X PUT --dump - http://localhost:8529/_api/replication/applier-start

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Stop replication applier

stop the replication

`PUT /_api/replication/applier-stop`

Stops the replication applier. This will return immediately if the replication applier is not running.

Example:

```
shell> curl -X PUT --dump - http://localhost:8529/_api/replication/applier-stop

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "state" : {
    "running" : false,
    "lastAppliedContinuousTick" : null,
    "lastProcessedContinuousTick" : null,
    "lastAvailableContinuousTick" : null,
    "safeResumeTick" : null,
    "progress" : {
      "time" : "2017-04-18T15:52:55Z",
      "message" : "fetching master state information",
      "failedConnects" : 1
    },
    "totalRequests" : 3,
    "totalFailedConnects" : 3,
    "totalEvents" : 0,
    "totalOperationsExcluded" : 0,
    "lastError" : {
      "errorNum" : 1412,
      "time" : "2017-04-18T15:52:55Z",
      "errorMessage" : "could not connect to master at tcp://127.0.0.1:8529: Could not connect to 'http+tcp://127.0.0
.1:8529' 'connect() failed with #111 - Connection refused'"
    },
    "time" : "2017-04-18T15:52:55Z"
  },
  "server" : {
    "version" : "3.1.18",
    "serverId" : "271514875687090"
  },
  "endpoint" : "tcp://127.0.0.1:8529",
  "database" : "_system"
}
```

Return Codes

- *200*: is returned if the request was executed successfully.
- *405*: is returned when an invalid HTTP method is used.
- *500*: is returned if an error occurred while assembling the response.

Examples

```
shell> curl -X PUT --dump - http://localhost:8529/_api/replication/applier-stop

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

State of the replication applier

output the current status of the replication

`GET /_api/replication/applier-state`

Returns the state of the replication applier, regardless of whether the applier is currently running or not.

The response is a JSON object with the following attributes:

- *state*: a JSON object with the following sub-attributes:
  - *running*: whether or not the applier is active and running
  - *lastAppliedContinuousTick*: the last tick value from the continuous replication log the applier has applied.
  - *lastProcessedContinuousTick*: the last tick value from the continuous replication log the applier has processed.

    Regularly, the last applied and last processed tick values should be identical. For transactional operations, the replication applier will first process incoming log events before applying them, so the processed tick value might be higher than the applied tick value. This will be the case until the applier encounters the *transaction commit* log event for the transaction.
  - *lastAvailableContinuousTick*: the last tick value the logger server can provide.
  - *time*: the time on the applier server.
  - *totalRequests*: the total number of requests the applier has made to the endpoint.
  - *totalFailedConnects*: the total number of failed connection attempts the applier has made.
  - *totalEvents*: the total number of log events the applier has processed.
  - *totalOperationsExcluded*: the total number of log events excluded because of *restrictCollections*.
  - *progress*: a JSON object with details about the replication applier progress. It contains the following sub-attributes if there is progress to report:
    - *message*: a textual description of the progress
    - *time*: the date and time the progress was logged
    - *failedConnects*: the current number of failed connection attempts
  - *lastError*: a JSON object with details about the last error that happened on the applier. It contains the following sub-attributes if there was an error:
    - *errorNum*: a numerical error code
    - *errorMessage*: a textual error description
    - *time*: the date and time the error occurred

    In case no error has occurred, *lastError* will be empty.
- *server*: a JSON object with the following sub-attributes:
  - *version*: the applier server's version
  - *serverId*: the applier server's id
- *endpoint*: the endpoint the applier is connected to (if applier is active) or will connect to (if applier is currently inactive)

- *database*: the name of the database the applier is connected to (if applier is active) or will connect to (if applier is currently inactive)

Example:

Fetching the state of an inactive applier:

```
shell> curl --dump - http://localhost:8529/_api/replication/applier-state

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "state" : {
    "running" : false,
    "lastAppliedContinuousTick" : null,
    "lastProcessedContinuousTick" : null,
    "lastAvailableContinuousTick" : null,
    "safeResumeTick" : null,
    "progress" : {
      "time" : "2017-04-18T15:52:54Z",
      "message" : "applier shut down",
      "failedConnects" : 1
    },
    "totalRequests" : 1,
    "totalFailedConnects" : 1,
    "totalEvents" : 0,
    "totalOperationsExcluded" : 0,
    "lastError" : {
      "errorNum" : 1412,
      "time" : "2017-04-18T15:52:54Z",
      "errorMessage" : "could not connect to master at tcp://127.0.0.1:8529: Could not connect to 'http+tcp://127.0.0
.1:8529' 'connect() failed with #111 - Connection refused'"
    },
    "time" : "2017-04-18T15:52:54Z"
  },
  "server" : {
    "version" : "3.1.18",
    "serverId" : "271514875687090"
  },
  "endpoint" : "tcp://127.0.0.1:8529",
  "database" : "_system"
}
```

Example:

Fetching the state of an active applier:

```
shell> curl --dump - http://localhost:8529/_api/replication/applier-state

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "state" : {
    "running" : true,
    "lastAppliedContinuousTick" : null,
    "lastProcessedContinuousTick" : null,
    "lastAvailableContinuousTick" : null,
    "safeResumeTick" : null,
    "progress" : {
      "time" : "2017-04-18T15:52:54Z",
      "message" : "applier shut down",
      "failedConnects" : 1
    },
    "totalRequests" : 1,
    "totalFailedConnects" : 1,
    "totalEvents" : 0,
    "totalOperationsExcluded" : 0,
    "lastError" : {
      "errorNum" : 0
    },
    "time" : "2017-04-18T15:52:54Z"
  },
  "server" : {
    "version" : "3.1.18",
    "serverId" : "271514875687090"
  },
  "endpoint" : "tcp://127.0.0.1:8529",
  "database" : "_system"
}
```

Return Codes

- *200*: is returned if the request was executed successfully.
- *405*: is returned when an invalid HTTP method is used.
- *500*: is returned if an error occurred while assembling the response.

Examples

Fetching the state of an inactive applier:

```
shell> curl --dump - http://localhost:8529/_api/replication/applier-state

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Fetching the state of an active applier:

```
shell> curl --dump - http://localhost:8529/_api/replication/applier-state

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Turn the server into a slave of another

Changes role to slave

```
PUT /_api/replication/make-slave
```

A json post document with these Properties is required:

- **username**: an optional ArangoDB username to use when connecting to the master.
- **includeSystem**: whether or not system collection operations will be applied
- **endpoint**: the master endpoint to connect to (e.g. "tcp://192.168.173.13:8529").
- **verbose**: if set to *true*, then a log line will be emitted for all operations performed by the replication applier. This should be used for debugging replication problems only.
- **connectTimeout**: the timeout (in seconds) when attempting to connect to the endpoint. This value is used for each connection attempt.
- **autoResync**: whether or not the slave should perform an automatic resynchronization with the master in case the master cannot serve log data requested by the slave, or when the replication is started and no tick value can be found.
- **database**: the database name on the master (if not specified, defaults to the name of the local current database).
- **idleMinWaitTime**: the minimum wait time (in seconds) that the applier will intentionally idle before fetching more log data from the master in case the master has already sent all its log data. This wait time can be used to control the frequency with which the replication applier sends HTTP log fetch requests to the master in case there is no write activity on the master. This value will be ignored if set to *0*.
- **requestTimeout**: the timeout (in seconds) for individual requests to the endpoint.
- **restrictType**: an optional string value for collection filtering. When specified, the allowed values are *include* or *exclude*.
- **idleMaxWaitTime**: the maximum wait time (in seconds) that the applier will intentionally idle before fetching more log data from the master in case the master has already sent all its log data and there have been previous log fetch attempts that resulted in no more log data. This wait time can be used to control the maximum frequency with which the replication applier sends HTTP log fetch requests to the master in case there is no write activity on the master for longer periods. This configuration value will only be used if the option *adaptivePolling* is set to *true*. This value will be ignored if set to *0*.
- **initialSyncMaxWaitTime**: the maximum wait time (in seconds) that the initial synchronization will wait for a response from the master when fetching initial collection data. This wait time can be used to control after what time the initial synchronization will give up waiting for a response and fail. This value is relevant even for continuous replication when *autoResync* is set to *true* because this may re-start the initial synchronization when the master cannot provide log data the slave requires. This value will be ignored if set to *0*.
- **restrictCollections** (string): an optional array of collections for use with *restrictType*. If *restrictType* is *include*, only the specified collections will be sychronised. If *restrictType* is *exclude*, all but the specified collections will be synchronized.
- **requireFromPresent**: if set to *true*, then the replication applier will check at start of its continuous replication if the start tick from the dump phase is still present on the master. If not, then there would be data loss. If *requireFromPresent* is *true*, the replication applier will abort with an appropriate error message. If set to *false*, then the replication applier will still start, and ignore the data loss.
- **adaptivePolling**: whether or not the replication applier will use adaptive polling.
- **maxConnectRetries**: the maximum number of connection attempts the applier will make in a row. If the applier cannot establish a connection to the endpoint in this number of attempts, it will stop itself.
- **password**: the password to use when connecting to the master.
- **connectionRetryWaitTime**: the time (in seconds) that the applier will intentionally idle before it retries connecting to the master in case of connection problems. This value will be ignored if set to *0*.
- **autoResyncRetries**: number of resynchronization retries that will be performed in a row when automatic resynchronization is enabled and kicks in. Setting this to *0* will effectively disable *autoResync*. Setting it to some other value will limit the number of retries that are performed. This helps preventing endless retries in case resynchronizations always fail.
- **chunkSize**: the requested maximum size for log transfer packets that is used when the endpoint is contacted.

Starts a full data synchronization from a remote endpoint into the local ArangoDB database and afterwards starts the continuous replication. The operation works on a per-database level.

All local database data will be removed prior to the synchronization.

In case of success, the body of the response is a JSON object with the following attributes:

- *state*: a JSON object with the following sub-attributes:
  - *running*: whether or not the applier is active and running
  - *lastAppliedContinuousTick*: the last tick value from the continuous replication log the applier has applied.
  - *lastProcessedContinuousTick*: the last tick value from the continuous replication log the applier has processed.

Regularly, the last applied and last processed tick values should be identical. For transactional operations, the replication applier will first process incoming log events before applying them, so the processed tick value might be higher than the applied tick value. This will be the case until the applier encounters the *transaction commit* log event for the transaction.

- *lastAvailableContinuousTick*: the last tick value the logger server can provide.
- *time*: the time on the applier server.
- *totalRequests*: the total number of requests the applier has made to the endpoint.
- *totalFailedConnects*: the total number of failed connection attempts the applier has made.
- *totalEvents*: the total number of log events the applier has processed.
- *totalOperationsExcluded*: the total number of log events excluded because of *restrictCollections*.
- *progress*: a JSON object with details about the replication applier progress. It contains the following sub-attributes if there is progress to report:
  - *message*: a textual description of the progress
  - *time*: the date and time the progress was logged
  - *failedConnects*: the current number of failed connection attempts
- *lastError*: a JSON object with details about the last error that happened on the applier. It contains the following sub-attributes if there was an error:
  - *errorNum*: a numerical error code
  - *errorMessage*: a textual error description
  - *time*: the date and time the error occurred

  In case no error has occurred, *lastError* will be empty.

- *server*: a JSON object with the following sub-attributes:
  - *version*: the applier server's version
  - *serverId*: the applier server's id
- *endpoint*: the endpoint the applier is connected to (if applier is active) or will connect to (if applier is currently inactive)
- *database*: the name of the database the applier is connected to (if applier is active) or will connect to (if applier is currently inactive)

WARNING: calling this method will sychronize data from the collections found on the remote master to the local ArangoDB database. All data in the local collections will be purged and replaced with data from the master.

Use with caution!

Please also keep in mind that this command may take a long time to complete and return. This is because it will first do a full data synchronization with the master, which will take time roughly proportional to the amount of data.

Note: this method is not supported on a coordinator in a cluster.

Return Codes

- *200*: is returned if the request was executed successfully.
- *400*: is returned if the configuration is incomplete or malformed.
- *405*: is returned when an invalid HTTP method is used.
- *500*: is returned if an error occurred during sychronization or when starting the continuous replication.
- *501*: is returned when this operation is called on a coordinator in a cluster.

Other Replication Commands

Return server id

fetch this server's unique identifier

`GET /_api/replication/server-id`

Returns the servers id. The id is also returned by other replication API methods, and this method is an easy means of determining a server's id.

The body of the response is a JSON object with the attribute *serverId*. The server id is returned as a string.

Example:

```
shell> curl --dump - http://localhost:8529/_api/replication/server-id

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "serverId" : "271514875687090"
}
```

Return Codes

- *200*: is returned if the request was executed successfully.
- *405*: is returned when an invalid HTTP method is used.
- *500*: is returned if an error occurred while assembling the response.

Examples

```
shell> curl --dump - http://localhost:8529/_api/replication/server-id

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "serverId" : "271514875687090"
}
```

HTTP Interface for Sharding

Sharding only should be used by developers!

Execute cluster roundtrip

executes a cluster roundtrip for sharding

`GET /_admin/cluster-test`

Executes a cluster roundtrip from a coordinator to a DB server and back. This call only works in a coordinator node in a cluster. One can and should append an arbitrary path to the URL and the part after */_admin/cluster-test* is used as the path of the HTTP request which is sent from the coordinator to a DB node. Likewise, any form data appended to the URL is forwarded in the request to the DB node. This handler takes care of all request types (see below) and uses the same request type in its request to the DB node.

The following HTTP headers are interpreted in a special way:

- *X-Shard-ID*: This specifies the ID of the shard to which the cluster request is sent and thus tells the system to which DB server to send the cluster request. Note that the mapping from the shard ID to the responsible server has to be defined in the agency under *Current/ShardLocation/*. One has to give this header, otherwise the system does not know where to send the request.
- *X-Client-Transaction-ID*: the value of this header is taken as the client transaction ID for the request
- *X-Timeout*: specifies a timeout in seconds for the cluster operation. If the answer does not arrive within the specified timeout, an corresponding error is returned and any subsequent real answer is ignored. The default if not given is 24 hours.
- *X-Synchronous-Mode*: If set to *true* the test function uses synchronous mode, otherwise the default asynchronous operation mode is used. This is mainly for debugging purposes.
- *Host*: This header is ignored and not forwarded to the DB server.
- *User-Agent*: This header is ignored and not forwarded to the DB server.

All other HTTP headers and the body of the request (if present, see other HTTP methods below) are forwarded as given in the original request.

In asynchronous mode the DB server answers with an HTTP request of its own, in synchronous mode it sends a HTTP response. In both cases the headers and the body are used to produce the HTTP response of this API call.

Return Codes

The return code can be anything the cluster request returns, as well as:

- *200*: is returned when everything went well, or if a timeout occurred. In the latter case a body of type application/json indicating the timeout is returned.
- *403*: is returned if ArangoDB is not running in cluster mode.
- *404*: is returned if ArangoDB was not compiled for cluster operation.

Execute cluster roundtrip

executes a cluster roundtrip for sharding

`POST /_admin/cluster-test`

Post Body

body (required): The body can be any type and is simply forwarded.

free style json body

See GET method.

Execute cluster roundtrip

executes a cluster roundtrip for sharding

`PUT /_admin/cluster-test`

Post Body

body (required):

free style json body

See GET method. The body can be any type and is simply forwarded.

Delete cluster roundtrip

executes a cluster roundtrip for sharding

`DELETE /_admin/cluster-test`

See GET method.

Update cluster roundtrip

executes a cluster roundtrip for sharding

`PATCH /_admin/cluster-test`

Post Body

body (required):

free style json body

See GET method. The body can be any type and is simply forwarded.

Execute cluster roundtrip

executes a cluster roundtrip for sharding

`HEAD /_admin/cluster-test`

See GET method.

Check port

allows to check whether a given port is usable

`GET /_admin/clusterCheckPort`

Query Parameters

- *port* (required):

Return Codes

- *200*: is returned when everything went well.
- *400*: the parameter port was not given or is no integer.

HTTP Interface for Administration and Monitoring

This is an introduction to ArangoDB's HTTP interface for administration and monitoring of the server.

Read global logs from the server

returns the server logs

`GET /_admin/log`

Query Parameters

- *upto* (optional): Returns all log entries up to log level *upto*. Note that *upto* must be:
- *fatal* or *0*
- *error* or *1*
- *warning* or *2*
- *info* or *3*
- *debug* or *4* The default value is *info*.
- *level* (optional): Returns all log entries of log level *level*. Note that the query parameters *upto* and *level* are mutually exclusive.
- *start* (optional): Returns all log entries such that their log entry identifier (*lid* value) is greater or equal to *start*.
- *size* (optional): Restricts the result to at most *size* log entries.
- *offset* (optional): Starts to return log entries skipping the first *offset* log entries. *offset* and *size* can be used for pagination.
- *search* (optional): Only return the log entries containing the text specified in *search*.
- *sort* (optional): Sort the log entries either ascending (if *sort* is *asc*) or descending (if *sort* is *desc*) according to their *lid* values. Note that the *lid* imposes a chronological order. The default value is *asc*.

Returns fatal, error, warning or info log messages from the server's global log. The result is a JSON object with the following attributes:

- *lid*: a list of log entry identifiers. Each log message is uniquely identified by its @LIT{lid} and the identifiers are in ascending order.
- *level*: a list of the log-levels for all log entries.
- *timestamp*: a list of the timestamps as seconds since 1970-01-01 for all log entries.
- *text* a list of the texts of all log entries
- *totalAmount*: the total amount of log entries before pagination.

Return Codes

- *400*: is returned if invalid values are specified for *upto* or *level*.
- *500*: is returned if the server cannot generate the result due to an out-of-memory error.

Return the current server loglevel

returns the current loglevel settings

`GET /_admin/log/level`

Returns the server's current loglevel settings. The result is a JSON object with the log topics being the object keys, and the log levels being the object values.

Return Codes

- *200*: is returned if the request is valid
- *500*: is returned if the server cannot generate the result due to an out-of-memory error.

Modify and return the current server loglevel

modifies the current loglevel settings

`PUT /_admin/log/level`

Modifies and returns the server's current loglevel settings. The request body must be a JSON object with the log topics being the object keys and the log levels being the object values.

The result is a JSON object with the adjusted log topics being the object keys, and the adjusted log levels being the object values.

Return Codes

- *200*: is returned if the request is valid
- *400*: is returned when the request body contains invalid JSON.
- *405*: is returned when an invalid HTTP method is used.
- *500*: is returned if the server cannot generate the result due to an out-of-memory error.

Reloads the routing information

Reload the routing table.

`POST /_admin/routing/reload`

Reloads the routing information from the collection *routing*.

Return Codes

- *200*: Routing information was reloaded successfully.

Read the statistics

return the statistics information

`GET /_admin/statistics`

Returns the statistics information. The returned object contains the statistics figures grouped together according to the description returned by *_admin/statistics-description*. For instance, to access a figure *userTime* from the group *system*, you first select the sub-object describing the group stored in *system* and in that sub-object the value for *userTime* is stored in the attribute of the same name.

In case of a distribution, the returned object contains the total count in *count* and the distribution list in *counts*. The sum (or total) of the individual values is returned in *sum*.

Example:

```
shell> curl --dump - http://localhost:8529/_admin/statistics

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "time" : 1492530751.123738,
  "enabled" : true,
  "system" : {
    "minorPageFaults" : 84657,
    "majorPageFaults" : 4,
    "userTime" : 3.74,
    "systemTime" : 1.28,
    "numberOfThreads" : 30,
    "residentSize" : 265711616,
    "residentSizePercent" : 0.0078776482977523,
    "virtualSize" : 4219047936
  },
  "client" : {
    "httpConnections" : 1,
    "connectionTime" : {
      "sum" : 0,
      "count" : 0,
      "counts" : [
        0,
        0,
        0,
        0
      ]
```

```
    },
    "totalTime" : {
      "sum" : 103.93337893486023,
      "count" : 3823,
      "counts" : [
        3578,
        124,
        102,
        15,
        1,
        0,
        3
      ]
    },
    "requestTime" : {
      "sum" : 60.11831212043762,
      "count" : 3823,
      "counts" : [
        3820,
        1,
        0,
        0,
        0,
        0,
        2
      ]
    },
    "queueTime" : {
      "sum" : 0,
      "count" : 0,
      "counts" : [
        0,
        0,
        0,
        0,
        0,
        0,
        0
      ]
    },
    "ioTime" : {
      "sum" : 43.81506681442261,
      "count" : 3823,
      "counts" : [
        3581,
        123,
        102,
        15,
        1,
        0,
        1
      ]
    },
    "bytesSent" : {
      "sum" : 1319525,
      "count" : 3823,
      "counts" : [
        427,
        3063,
        327,
        6,
        0,
        0
      ]
    },
    "bytesReceived" : {
      "sum" : 828697,
      "count" : 3823,
      "counts" : [
        3536,
        287,
        0,
        0,
        0,
        0
      ]
```

```
    }
  },
  "http" : {
    "requestsTotal" : 3823,
    "requestsAsync" : 11,
    "requestsGet" : 811,
    "requestsHead" : 0,
    "requestsPost" : 2794,
    "requestsPut" : 66,
    "requestsPatch" : 3,
    "requestsDelete" : 149,
    "requestsOptions" : 0,
    "requestsOther" : 0
  },
  "server" : {
    "uptime" : 50.755594968795776,
    "physicalMemory" : 33729814528
  },
  "error" : false,
  "code" : 200
}
```

Return Codes

- *200*: Statistics were returned successfully.

Examples

```
shell> curl --dump - http://localhost:8529/_admin/statistics

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Statistics description

fetch descriptive info of statistics

`GET /_admin/statistics-description`

Returns a description of the statistics returned by */_admin/statistics*. The returned objects contains an array of statistics groups in the attribute *groups* and an array of statistics figures in the attribute *figures*.

A statistics group is described by

- *group*: The identifier of the group.
- *name*: The name of the group.
- *description*: A description of the group.

A statistics figure is described by

- *group*: The identifier of the group to which this figure belongs.
- *identifier*: The identifier of the figure. It is unique within the group.
- *name*: The name of the figure.
- *description*: A description of the figure.
- *type*: Either *current*, *accumulated*, or *distribution*.
- *cuts*: The distribution vector.
- *units*: Units in which the figure is measured.

Example:

```
shell> curl --dump - http://localhost:8529/_admin/statistics-description
```

```
HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "groups" : [
    {
      "group" : "system",
      "name" : "Process Statistics",
      "description" : "Statistics about the ArangoDB process"
    },
    {
      "group" : "client",
      "name" : "Client Connection Statistics",
      "description" : "Statistics about the connections."
    },
    {
      "group" : "http",
      "name" : "HTTP Request Statistics",
      "description" : "Statistics about the HTTP requests."
    },
    {
      "group" : "server",
      "name" : "Server Statistics",
      "description" : "Statistics about the ArangoDB server"
    }
  ],
  "figures" : [
    {
      "group" : "system",
      "identifier" : "userTime",
      "name" : "User Time",
      "description" : "Amount of time that this process has been scheduled in user mode, measured in seconds.",
      "type" : "accumulated",
      "units" : "seconds"
    },
    {
      "group" : "system",
      "identifier" : "systemTime",
      "name" : "System Time",
      "description" : "Amount of time that this process has been scheduled in kernel mode, measured in seconds.",
      "type" : "accumulated",
      "units" : "seconds"
    },
    {
      "group" : "system",
      "identifier" : "numberOfThreads",
      "name" : "Number of Threads",
      "description" : "Number of threads in the arangod process.",
      "type" : "current",
      "units" : "number"
    },
    {
      "group" : "system",
      "identifier" : "residentSize",
      "name" : "Resident Set Size",
      "description" : "The total size of the number of pages the process has in real memory. This is just the pages w
hich count toward text, data, or stack space. This does not include pages which have not been demand-loaded in, or wh
ich are swapped out. The resident set size is reported in bytes.",
      "type" : "current",
      "units" : "bytes"
    },
    {
      "group" : "system",
      "identifier" : "residentSizePercent",
      "name" : "Resident Set Size",
      "description" : "The percentage of physical memory used by the process as resident set size.",
      "type" : "current",
      "units" : "percent"
    },
    {
      "group" : "system",
      "identifier" : "virtualSize",
      "name" : "Virtual Memory Size",
      "description" : "On Windows, this figure contains the total amount of memory that the memory manager has commit
ted for the arangod process. On other systems, this figure contains The size of the virtual memory the process is usi
```

```
ng.",
      "type" : "current",
      "units" : "bytes"
    },
    {
      "group" : "system",
      "identifier" : "minorPageFaults",
      "name" : "Minor Page Faults",
      "description" : "The number of minor faults the process has made which have not required loading a memory page
from disk. This figure is not reported on Windows.",
      "type" : "accumulated",
      "units" : "number"
    },
    {
      "group" : "system",
      "identifier" : "majorPageFaults",
      "name" : "Major Page Faults",
      "description" : "On Windows, this figure contains the total number of page faults. On other system, this figure
 contains the number of major faults the process has made which have required loading a memory page from disk.",
      "type" : "accumulated",
      "units" : "number"
    },
    {
      "group" : "client",
      "identifier" : "httpConnections",
      "name" : "Client Connections",
      "description" : "The number of connections that are currently open.",
      "type" : "current",
      "units" : "number"
    },
    {
      "group" : "client",
      "identifier" : "totalTime",
      "name" : "Total Time",
      "description" : "Total time needed to answer a request.",
      "type" : "distribution",
      "cuts" : [
        0.01,
        0.05,
        0.1,
        0.2,
        0.5,
        1
      ],
      "units" : "seconds"
    },
    {
      "group" : "client",
      "identifier" : "requestTime",
      "name" : "Request Time",
      "description" : "Request time needed to answer a request.",
      "type" : "distribution",
      "cuts" : [
        0.01,
        0.05,
        0.1,
        0.2,
        0.5,
        1
      ],
      "units" : "seconds"
    },
    {
      "group" : "client",
      "identifier" : "queueTime",
      "name" : "Queue Time",
      "description" : "Queue time needed to answer a request.",
      "type" : "distribution",
      "cuts" : [
        0.01,
        0.05,
        0.1,
        0.2,
        0.5,
        1
      ],
```

```
      "units" : "seconds"
    },
    {
      "group" : "client",
      "identifier" : "bytesSent",
      "name" : "Bytes Sent",
      "description" : "Bytes sents for a request.",
      "type" : "distribution",
      "cuts" : [
        250,
        1000,
        2000,
        5000,
        10000
      ],
      "units" : "bytes"
    },
    {
      "group" : "client",
      "identifier" : "bytesReceived",
      "name" : "Bytes Received",
      "description" : "Bytes receiveds for a request.",
      "type" : "distribution",
      "cuts" : [
        250,
        1000,
        2000,
        5000,
        10000
      ],
      "units" : "bytes"
    },
    {
      "group" : "client",
      "identifier" : "connectionTime",
      "name" : "Connection Time",
      "description" : "Total connection time of a client.",
      "type" : "distribution",
      "cuts" : [
        0.1,
        1,
        60
      ],
      "units" : "seconds"
    },
    {
      "group" : "http",
      "identifier" : "requestsTotal",
      "name" : "Total requests",
      "description" : "Total number of HTTP requests.",
      "type" : "accumulated",
      "units" : "number"
    },
    {
      "group" : "http",
      "identifier" : "requestsAsync",
      "name" : "Async requests",
      "description" : "Number of asynchronously executed HTTP requests.",
      "type" : "accumulated",
      "units" : "number"
    },
    {
      "group" : "http",
      "identifier" : "requestsGet",
      "name" : "HTTP GET requests",
      "description" : "Number of HTTP GET requests.",
      "type" : "accumulated",
      "units" : "number"
    },
    {
      "group" : "http",
      "identifier" : "requestsHead",
      "name" : "HTTP HEAD requests",
      "description" : "Number of HTTP HEAD requests.",
      "type" : "accumulated",
      "units" : "number"
```

```
    },
    {
      "group" : "http",
      "identifier" : "requestsPost",
      "name" : "HTTP POST requests",
      "description" : "Number of HTTP POST requests.",
      "type" : "accumulated",
      "units" : "number"
    },
    {
      "group" : "http",
      "identifier" : "requestsPut",
      "name" : "HTTP PUT requests",
      "description" : "Number of HTTP PUT requests.",
      "type" : "accumulated",
      "units" : "number"
    },
    {
      "group" : "http",
      "identifier" : "requestsPatch",
      "name" : "HTTP PATCH requests",
      "description" : "Number of HTTP PATCH requests.",
      "type" : "accumulated",
      "units" : "number"
    },
    {
      "group" : "http",
      "identifier" : "requestsDelete",
      "name" : "HTTP DELETE requests",
      "description" : "Number of HTTP DELETE requests.",
      "type" : "accumulated",
      "units" : "number"
    },
    {
      "group" : "http",
      "identifier" : "requestsOptions",
      "name" : "HTTP OPTIONS requests",
      "description" : "Number of HTTP OPTIONS requests.",
      "type" : "accumulated",
      "units" : "number"
    },
    {
      "group" : "http",
      "identifier" : "requestsOther",
      "name" : "other HTTP requests",
      "description" : "Number of other HTTP requests.",
      "type" : "accumulated",
      "units" : "number"
    },
    {
      "group" : "server",
      "identifier" : "uptime",
      "name" : "Server Uptime",
      "description" : "Number of seconds elapsed since server start.",
      "type" : "current",
      "units" : "seconds"
    },
    {
      "group" : "server",
      "identifier" : "physicalMemory",
      "name" : "Physical Memory",
      "description" : "Physical memory in bytes.",
      "type" : "current",
      "units" : "bytes"
    }
  ],
  "error" : false,
  "code" : 200
}
```

Return Codes

- *200*: Description was returned successfully.

Examples

```
shell> curl --dump - http://localhost:8529/_admin/statistics-description

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Return role of a server in a cluster

Get to know whether this server is a Coordinator or DB-Server

`GET /_admin/server/role`

Returns the role of a server in a cluster. The role is returned in the *role* attribute of the result. Possible return values for *role* are:

- *SINGLE*: the server is a standalone server without clustering
- *COORDINATOR*: the server is a coordinator in a cluster
- *PRIMARY*: the server is a primary database server in a cluster
- *SECONDARY*: the server is a secondary database server in a cluster
- *AGENT*: the server is an agency node in a cluster
- *UNDEFINED*: in a cluster, *UNDEFINED* is returned if the server role cannot be determined.

Return Codes

- *200*: Is returned in all cases.

Return id of a server in a cluster

Get to know the internal id of the server

`GET /_admin/server/id`

Returns the id of a server in a cluster. The id is empty if it is working in single server mode.

Return Codes

- *200*: Is returned in all cases.

Queries statistics of DBserver

allows to query the statistics of a DBserver in the cluster

`GET /_admin/clusterStatistics`

Query Parameters

- *DBserver* (required):

Return Codes

- *200*: is returned when everything went well.
- *400*: the parameter DBserver was not given or is not the ID of a DBserver
- *403*: server is not a coordinator.

HTTP Interface for Endpoints

The API `/_api/endpoint` is *deprecated*. For cluster mode there is `/_api/cluster/endpoints` to find all current coordinator endpoints (see below).

The ArangoDB server can listen for incoming requests on multiple *endpoints*.

The endpoints are normally specified either in ArangoDB's configuration file or on the command-line, using the "--server.endpoint" option. The default endpoint for ArangoDB is *tcp://127.0.0.1:8529* or *tcp://localhost:8529*.

Please note that all endpoint management operations can only be accessed via the default database (*_system*) and none of the other databases.

Asking about Endpoints via HTTP

Return list of all endpoints

This API call returns the list of all endpoints (single server).

`GET /_api/endpoint`

THIS API IS DEPRECATED

Returns an array of all configured endpoints the server is listening on.

The result is a JSON array of JSON objects, each with `"entrypoint"' as the only attribute, and with the value being a string describing the endpoint.

Note: retrieving the array of all endpoints is allowed in the system database only. Calling this action in any other database will make the server return an error.

Example:

```
shell> curl --dump - http://localhost:8529/_api/endpoint

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

[
  {
    "endpoint" : "http://127.0.0.1:60108"
  }
]
```

Return Codes

- *200*: is returned when the array of endpoints can be determined successfully.
- *400*: is returned if the action is not carried out in the system database.
- *405*: The server will respond with *HTTP 405* if an unsupported HTTP method is used.

Examples

```
shell> curl --dump - http://localhost:8529/_api/endpoint

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

[
  {
    "endpoint" : "http://127.0.0.1:60108"
  }
]
```

Get information about all coordinator endpoints

This API call returns information about all coordinator endpoints (cluster only).

`GET /_api/cluster/endpoints`

Returns an object with an attribute `endpoints`, which contains an array of objects, which each have the attribute `endpoint`, whose value is a string with the endpoint description. There is an entry for each coordinator in the cluster. This method only works on coordinators in cluster mode. In case of an error the `error` attribute is set to `true`.

Return Codes

- *200*: is returned when everything went well.
- *403*: server is not a coordinator or method was not GET.

HTTP Interface for User Management

This is an introduction to ArangoDB's HTTP interface for managing users.

The interface provides a simple means to add, update, and remove users. All users managed through this interface will be stored in the system collection *_users*. You should never manipulate the *_users* collection directly.

This specialized interface intentionally does not provide all functionality that is available in the regular document REST API.

Please note that user operations are not included in ArangoDB's replication.

Create User

Create a new user.

`POST /_api/user`

A json post document with these Properties is required:

- **passwd**: The user password as a string. If no password is specified, the empty string will be used. If you pass the special value *ARANGODB_DEFAULT_ROOT_PASSWORD*, the password will be set the value stored in the environment variable `ARANGODB_DEFAULT_ROOT_PASSWORD`. This can be used to pass an instance variable into ArangoDB. For example, the instance identifier from Amazon.
- **active**: An optional flag that specifies whether the user is active. If not specified, this will default to true
- **user**: The name of the user as a string. This is mandatory.
- **extra**: An optional JSON object with arbitrary extra data about the user.

Create a new user. This user will not have access to any database. You need permission to the *_system* database in order to execute this REST call.

Example:

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/user <<EOF
{
  "user" : "admin@example",
  "passwd" : "secure"
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8

{
  "user" : "admin@example",
  "active" : true,
  "extra" : {
  },
  "changePassword" : false,
  "error" : false,
  "code" : 201
}
```

Return Codes

- *201*: Returned if the user can be added by the server
- *400*: If the JSON representation is malformed or mandatory data is missing from the request.
- *409*: Returned if a user with the same name already exists.

Examples

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/user <<EOF
{
  "user" : "admin@example",
  "passwd" : "secure"
}
EOF

HTTP/1.1 201 Created
content-type: application/json; charset=utf-8
```

show response body

Grant or revoke database access

Grant or revoke access to a database.

`PUT /_api/user/{user}/database/{dbname}`

A json post document with these Properties is required:

- **grant**: Use "rw" to grant access right and "none" to revoke.

Grants or revokes access to the database *dbname* for user *user*. You need permission to the *_system* database in order to execute this REST call.

Example:

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/user/admin@myapp/database/_system <<EOF
{
  "grant" : "rw"
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "_system" : "rw",
  "error" : false,
  "code" : 200
}
```

Path Parameters

- *user* (required): The name of the user.
- *dbname* (required): The name of the database.

Return Codes

- *200*: Returned if the access permissions were changed successfully.
- *400*: If the JSON representation is malformed or mandatory data is missing from the request.

Examples

```
shell> curl -X PUT --data-binary @- --dump -
http://localhost:8529/_api/user/admin@myapp/database/_system <<EOF
{
  "grant" : "rw"
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

List the databases available to a User

List available database to the specified user

`GET /_api/user/{user}/database/`

Path Parameters

- *user* (required): The name of the user for which you want to query the databases.

Fetch the list of databases available to the specified *user*. You need permission to the *_system* database in order to execute this REST call.

The call will return a JSON object with the per-database access privileges for the specified user. The *result* object will contain the databases names as object keys, and the associated privileges for the database as values.

Example:

```
shell> curl --dump - http://localhost:8529/_api/user/anotherAdmin@secapp/database/

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "result" : {
    "_system" : "rw"
  },
  "error" : false,
  "code" : 200
}
```

Return Codes

- *200*: Returned if the list of available databases can be returned.
- *400*: If the access privileges are not right etc.

Examples

```
shell> curl --dump - http://localhost:8529/_api/user/anotherAdmin@secapp/database/

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Replace User

replace an existing user with a new one.

`PUT /_api/user/{user}`

Path Parameters

- *user* (required): The name of the user

A json post document with these Properties is required:

- **passwd**: The user password as a string. Specifying a password is mandatory, but the empty string is allowed for passwords
- **active**: An optional flag that specifies whether the user is active. If not specified, this will default to true
- **extra**: An optional JSON object with arbitrary extra data about the user.

Replaces the data of an existing user. The name of an existing user must be specified in user. You can only change the password of your self. You need access to the *_system* database to change the *active* flag.

Example:

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/user/admin@myapp <<EOF
{
  "passwd" : "secure"
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "user" : "admin@myapp",
  "active" : true,
  "extra" : {
  },
  "changePassword" : false,
  "error" : false,
  "code" : 200
}
```

Return Codes

- *200*: Is returned if the user data can be replaced by the server
- *400*: The JSON representation is malformed or mandatory data is missing from the request
- *404*: The specified user does not exist

Examples

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/user/admin@myapp
<<EOF
{
  "passwd" : "secure"
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Update User

modify attributes of an existing user

`PATCH /_api/user/{user}`

Path Parameters

- *user* (required): The name of the user

A json post document with these Properties is required:

- **passwd**: The user password as a string. Specifying a password is mandatory, but the empty string is allowed for passwords
- **active**: An optional flag that specifies whether the user is active. If not specified, this will default to true
- **extra**: An optional JSON object with arbitrary extra data about the user.

Partially updates the data of an existing user. The name of an existing user must be specified in *user*. You can only change the password of your self. You need access to the *_system* database to change the *active* flag.

Example:

```
shell> curl -X PATCH --data-binary @- --dump - http://localhost:8529/_api/user/admin@myapp <<EOF
{
  "passwd" : "secure"
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "user" : "admin@myapp",
  "active" : true,
  "extra" : {
  },
  "changePassword" : false,
  "error" : false,
  "code" : 200
}
```

Return Codes

- *200*: Is returned if the user data can be replaced by the server
- *400*: The JSON representation is malformed or mandatory data is missing from the request
- *404*: The specified user does not exist

Examples

```
shell> curl -X PATCH --data-binary @- --dump -
http://localhost:8529/_api/user/admin@myapp <<EOF
{
  "passwd" : "secure"
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Remove User

delete a user permanently.

`DELETE /_api/user/{user}`

Path Parameters

- *user* (required): The name of the user

Removes an existing user, identified by *user*. You need access to the *_system* database.

Example:

```
shell> curl -X DELETE --data-binary @- --dump - http://localhost:8529/_api/user/userToDelete@myapp <<EOF
{
}
EOF

HTTP/1.1 202 Accepted
content-type: application/json; charset=utf-8

{
  "error" : false,
  "code" : 202
}
```

Return Codes

- *202*: Is returned if the user was removed by the server
- *404*: The specified user does not exist

Examples

```
shell> curl -X DELETE --data-binary @- --dump -
http://localhost:8529/_api/user/userToDelete@myapp <<EOF
{
}
EOF

HTTP/1.1 202 Accepted
content-type: application/json; charset=utf-8

{
  "error" : false,
  "code" : 202
}
```

Fetch User

fetch the properties of a user.

`GET /_api/user/{user}`

Path Parameters

- *user* (required): The name of the user

Fetches data about the specified user. You can fetch information about yourself or you need permission to the *_system* database in order to execute this REST call.

Example:

```
shell> curl --dump - http://localhost:8529/_api/user/admin@myapp

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "user" : "admin@myapp",
  "active" : true,
  "extra" : {
  },
  "changePassword" : false,
  "error" : false,
  "code" : 200
}
```

Return Codes

- *200*: The user was found.
- *404*: The user with the specified name does not exist.

Examples

```
shell> curl --dump - http://localhost:8529/_api/user/admin@myapp

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

List available Users

fetch the properties of a user.

`GET /_api/user/`

Fetches data about all users. You can only execute this call if you have access to the *_system* database. Otherwise, you will only get information about yourself.

The call will return a JSON object with at least the following attributes on success:

- *user*: The name of the user as a string.
- *active*: An optional flag that specifies whether the user is active.
- *extra*: An optional JSON object with arbitrary extra data about the user.
- *changePassword*: An optional flag that specifies whether the user must change the password or not.

Example:

```
shell> curl --dump - http://localhost:8529/_api/user

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "result" : [
    {
      "user" : "root",
      "active" : true,
      "extra" : {
      },
      "changePassword" : false
    },
    {
      "user" : "admin",
      "active" : true,
      "extra" : {
      },
      "changePassword" : false
    },
    {
      "user" : "tester",
      "active" : false,
      "extra" : {
      },
      "changePassword" : false
    },
    {
      "user" : "admin@example",
      "active" : true,
      "extra" : {
      },
      "changePassword" : false
    }
  ],
  "error" : false,
  "code" : 200
}
```

Return Codes

- *200*: The users that were found.

Examples

```
shell> curl --dump - http://localhost:8529/_api/user

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

HTTP tasks Interface

Following you have ArangoDB's HTTP Interface for Tasks.

There are also some examples provided for every API action.

Fetch all tasks or one task

Retrieves all currently active server tasks

`GET /_api/tasks/`

fetches all existing tasks on the server

Example:

Fetching all tasks

```
shell> curl --dump - http://localhost:8529/_api/tasks

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

[
  {
    "id" : "statistics-gc",
    "name" : "statistics-gc",
    "created" : 1492530712.579428,
    "type" : "periodic",
    "period" : 450,
    "offset" : 72.262448,
    "command" : "(function (params) { require('@arangodb/statistics').garbageCollector(); } )(params);",
    "database" : "_system"
  },
  {
    "id" : "92",
    "name" : "user-defined task",
    "created" : 1492530702.831905,
    "type" : "periodic",
    "period" : 1,
    "offset" : 0.000001,
    "command" : "(function (params) { (function () {
        require('@arangodb/foxx/queues/manager').manage();
      })(params) } )(params);",
    "database" : "_system"
  },
  {
    "id" : "statistics-average-collector",
    "name" : "statistics-average-collector",
    "created" : 1492530712.579171,
    "type" : "periodic",
    "period" : 900,
    "offset" : 20,
    "command" : "(function (params) { require('@arangodb/statistics').historianAverage(); } )(params);",
    "database" : "_system"
  },
  {
    "id" : "statistics-collector",
    "name" : "statistics-collector",
    "created" : 1492530712.57898,
    "type" : "periodic",
    "period" : 10,
    "offset" : 1,
    "command" : "(function (params) { require('@arangodb/statistics').historian(); } )(params);",
    "database" : "_system"
  }
]
```

- *200*: The list of tasks

Examples

Fetching all tasks

```
shell> curl --dump - http://localhost:8529/_api/tasks

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

[
  {
    "id" : "statistics-gc",
    "name" : "statistics-gc",
    "created" : 1492530712.579428,
    "type" : "periodic",
    "period" : 450,
    "offset" : 72.262448,
    "command" : "(function (params) {
require('@arangodb/statistics').garbageCollector(); } )(params);",
    "database" : "_system"
  },
  {
    "id" : "92",
    "name" : "user-defined task",
    "created" : 1492530702.831905,
    "type" : "periodic",
    "period" : 1,
    "offset" : 0.000001,
    "command" : "(function (params) { (function () {\n
require('@arangodb/foxx/queues/manager').manage();\n      })(params) } )(params);",
    "database" : "_system"
  },
  {
    "id" : "statistics-average-collector",
    "name" : "statistics-average-collector",
    "created" : 1492530712.579171,
    "type" : "periodic",
    "period" : 900,
    "offset" : 20,
    "command" : "(function (params) {
require('@arangodb/statistics').historianAverage(); } )(params);",
    "database" : "_system"
  },
  {
    "id" : "statistics-collector",
    "name" : "statistics-collector",
    "created" : 1492530712.57898,
    "type" : "periodic",
    "period" : 10,
    "offset" : 1,
    "command" : "(function (params) { require('@arangodb/statistics').historian(); } )
(params);",
    "database" : "_system"
  }
]
```

Fetch one task with id

Retrieves one currently active server task

`GET /_api/tasks/{id}`

- *id* (required): The id of the task to fetch.

fetches one existing tasks on the server specified by *id*

Example:

Fetching a single task by its id

```
shell> curl --dump - http://localhost:8529/_api/tasks/statistics-average-collector

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "id" : "statistics-average-collector",
  "name" : "statistics-average-collector",
  "created" : 1492530712.579171,
  "type" : "periodic",
  "period" : 900,
  "offset" : 20,
  "command" : "(function (params) { require('@arangodb/statistics').historianAverage(); } )(params);",
  "database" : "_system",
  "error" : false,
  "code" : 200
}
```

Example:

trying to fetch a non-existing task

```
shell> curl --dump - http://localhost:8529/_api/tasks/non-existing-task

HTTP/1.1 404 Not Found
content-type: application/json; charset=utf-8

{
  "error" : true,
  "code" : 404,
  "errorNum" : 1852,
  "errorMessage" : "task not found"
}
```

- *200*: The requested task

Examples

Fetching a single task by its id

```
shell> curl --dump - http://localhost:8529/_api/tasks/statistics-average-collector

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body
trying to fetch a non-existing task

```
shell> curl --dump - http://localhost:8529/_api/tasks/non-existing-task

HTTP/1.1 404 Not Found
content-type: application/json; charset=utf-8
```

show response body

creates a task

creates a new task

`POST /_api/tasks`

A json post document with these Properties is required:

- **params**: The parameters to be passed into command
- **offset**: Number of seconds initial delay
- **command**: The JavaScript code to be executed
- **name**: The name of the task
- **period**: number of seconds between the executions

creates a new task with a generated id

Example:

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/tasks/ <<EOF
{
  "name" : "SampleTask",
  "command" : "(function(params) { require('@arangodb').print(params); })(params)",
  "params" : {
    "foo" : "bar",
    "bar" : "foo"
  },
  "period" : 2
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "id" : "11270",
  "name" : "SampleTask",
  "created" : 1492530804.323003,
  "type" : "periodic",
  "period" : 2,
  "offset" : 0,
  "command" : "(function (params) { (function(params) { require('@arangodb').print(params); })(params) } )(params);",

  "database" : "_system",
  "error" : false,
  "code" : 200
}
shell> curl -X DELETE --dump - http://localhost:8529/_api/tasks/11270
```

Return Codes

- *400*: If the post body is not accurate, a *HTTP 400* is returned.

Examples

```
shell> curl -X POST --data-binary @- --dump - http://localhost:8529/_api/tasks/ <<EOF
{
  "name" : "SampleTask",
  "command" : "(function(params) { require('@arangodb').print(params); })(params)",
  "params" : {
    "foo" : "bar",
    "bar" : "foo"
  },
  "period" : 2
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

creates a task with id

registers a new task with a pre-defined id

`PUT /_api/tasks/{id}`

- *id* (required): The id of the task to create

A json post document with these Properties is required:

- **params**: The parameters to be passed into command
- **offset**: Number of seconds initial delay
- **command**: The JavaScript code to be executed
- **name**: The name of the task
- **period**: number of seconds between the executions

registers a new task with the specified id

Example:

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/tasks/sampleTask <<EOF
{
  "id" : "SampleTask",
  "name" : "SampleTask",
  "command" : "(function(params) { require('@arangodb').print(params); })(params)",
  "params" : {
    "foo" : "bar",
    "bar" : "foo"
  },
  "period" : 2
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "id" : "sampleTask",
  "name" : "SampleTask",
  "created" : 1492530804.357697,
  "type" : "periodic",
  "period" : 2,
  "offset" : 0,
  "command" : "(function (params) { (function(params) { require('@arangodb').print(params); })(params) } )(params);",

  "database" : "_system",
  "error" : false,
  "code" : 200
}
```

Return Codes

- *400*: If the task *id* already exists or the rest body is not accurate, *HTTP 400* is returned.

Examples

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_api/tasks/sampleTask
<<EOF
{
  "id" : "SampleTask",
  "name" : "SampleTask",
  "command" : "(function(params) { require('@arangodb').print(params); })(params)",
  "params" : {
    "foo" : "bar",
    "bar" : "foo"
  },
  "period" : 2
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

deletes the task with id

deletes one currently active server task

`DELETE /_api/tasks/{id}`

- *id* (required): The id of the task to delete.

Deletes the task identified by *id* on the server.

Example:

trying to delete non existing task

```
shell> curl -X DELETE --dump - http://localhost:8529/_api/tasks/NoTaskWithThatName

HTTP/1.1 404 Not Found
content-type: application/json; charset=utf-8

{
  "error" : true,
  "code" : 404,
  "errorNum" : 1852,
  "errorMessage" : "task not found"
}
```

Example:

Remove existing Task

```
shell> curl -X DELETE --dump - http://localhost:8529/_api/tasks/SampleTask

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "error" : false,
  "code" : 200
}
```

Return Codes

- *404*: If the task *id* is unknown, then an *HTTP 404* is returned.

Examples

trying to delete non existing task

```
shell> curl -X DELETE --dump - http://localhost:8529/_api/tasks/NoTaskWithThatName

HTTP/1.1 404 Not Found
content-type: application/json; charset=utf-8
```

show response body
Remove existing Task

```
shell> curl -X DELETE --dump - http://localhost:8529/_api/tasks/SampleTask

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "error" : false,
  "code" : 200
}
```

HTTP Interface for Agency feature

Configuration

At all times, i.e. regardless of the state of the agents and the current health of the RAFT consensus, one can invoke the configuration API:

```
http://server:port/_api/agency/config
```

The output might look somewhat like this

```
{
  "term": 1,
  "leaderId": "f5d11cde-8468-4fd2-8747-b4ef5c7dfa98",
  "lastCommitted": 1,
  "lastAcked": {
    "ac129027-b440-4c4f-84e9-75c042942171": 0.21,
    "c54dbb8a-723d-4c82-98de-8c841a14a112": 0.21,
    "f5d11cde-8468-4fd2-8747-b4ef5c7dfa98": 0
  },
  "configuration": {
    "pool": {
      "ac129027-b440-4c4f-84e9-75c042942171": "tcp://localhost:8531",
      "c54dbb8a-723d-4c82-98de-8c841a14a112": "tcp://localhost:8530",
      "f5d11cde-8468-4fd2-8747-b4ef5c7dfa98": "tcp://localhost:8529"
    },
    "active": [
      "ac129027-b440-4c4f-84e9-75c042942171",
      "c54dbb8a-723d-4c82-98de-8c841a14a112",
      "f5d11cde-8468-4fd2-8747-b4ef5c7dfa98"
    ],
    "id": "f5d11cde-8468-4fd2-8747-b4ef5c7dfa98",
    "agency size": 3,
    "pool size": 3,
    "endpoint": "tcp://localhost:8529",
    "min ping": 0.5,
    "max ping": 2.5,
    "supervision": false,
    "supervision frequency": 5,
    "compaction step size": 1000,
    "supervision grace period": 120
  }
}
```

This is the actual output of a healthy agency. The configuration of the agency is found in the `configuration` section as you might have guessed. It is populated by static information on the startup parameters like `agency size`, the once generated `unique id` etc. It holds information on the invariants of the RAFT algorithm and data compaction.

The remaining data reflect the variant entities in RAFT, as `term` and `leaderId`, also some debug information on how long the last leadership vote was received from any particular agency member. Low term numbers on a healthy network are an indication of good operation environemnt, while often increasing term numbers indicate, that the network environemnt and stability suggest to raise the RAFT parameters `min ping` and 'max ping' accordingly.

Key-Value store APIs

Generally, all document IO to and from the key-value store consists of JSON arrays. The outer Array is an envelope for multiple read or write transactions. The results are arrays are an envelope around the results corresponding to the order of the incoming transactions.

Consider the following write operation into a prestine agency:

```
curl -L http://server:port/write -d '[[{"a":{"b":{"c":[1,2,3]},"e":12},"d":false}]]'
```

```
[{results:[1]}]
```

And the subsequent read operation

```
curl -L http://server:port/read -d '[["/"]]'
```

```
[
  {
    "a": {
      "b": {
        "c": [1,2,3]
      },
      "e": 12
    },
    "d": false
  }
]
```

In the first step we commited a single transaction that commits the JSON document inside the inner transaction array to the agency. The result is `[1]`, which is the replicated log index. Repeated invocation will yield growing log numbers 2, 3, 4, etc.

The read access is a complete access to the key-value store indicated by access to it's root element and returns the result as an array corresponding to the outermost array in the read transaction.

Let's dig in some deeper.

Read API

Let's start with the above initialised key-value store in the following. Let us visit the following read operations:

```
curl -L http://server:port/read -d '[["/a/b"]]'
```

```
[
  {
    "a": {
      "b": {
        "c": [1,2,3]
      }
    }
  }
]
```

And

```
curl -L http://server:port/read -d '[["/a/b/c"]]'
```

```
[
  {
    "a": {
      "b": {
        "c": [1,2,3]
      }
    }
  }
]
```

Note that the above results are identical, meaning that results obtained from the agencyare always return with full path.

The second outer array brackets in read operations correspond to transactions, meaning that the result is guaranteed to have been acquired without a write transaction in between:

```
curl -L http://server:port/read -d '[["/a/e"],["/d","/a/b"]]'
```

```
[
  {
    "a": {
      "e": 12
    }
  },
  {
    "a": {
      "b": {
        "c": [1,2,3
        ]
      }
    },
    "d": false
  }
]
```

While the first transaction consists of a single read access to the key-value-store thus strechting the meaning of the word transaction, the second bracket actually hold two disjunct read accesses, which have been joined within zero-time, i.e. without a write access in between. That is to say that `"/d"` cannot have changed before `"/a/b"` had been acquired.

Let's try to fetch a value from the key-value-store, which does not exist:

```
curl -L http://server:port/read -d '[["/a/b/d"]]'
```

```
[
  {
    "a": {
      "b": {}
    }
  }
]
```

The result returns the cross section of the requested path and the key-value-store contents. `"/a/b"` exists, but there is no key `"/a/b/d"`. Thus the following transaction will yield:

```
curl -L http://server:port/read -d '[["/a/b/d","/d"]]'
```

```
[
  {
    "a": {
      "b": {}
    },
    "d": false
  }
]
```

And this last read operation should return:

```
curl -L http://server:port/read -d '[["/a/b/c"],["/a/b/d"],["/a/x/y"],["/y"],["/a/b","/a/x" ]]'
```

```
[
  {"a":{"b":{"c":[1,2,3]}}},
  {"a":{"b":{}}},
  {"a":{}},
  {},
  {"a":{"b":{"c":[1,2,3]}}}
]
```

Write API

The write API must obviously be more versatile and needs a more detailed appreciation. Write operations are arrays of transactions with preconditions, i.e. `[[U,P]]`, where the system tries to apply all updates in the outer array in turn, rejecting those whose precondition is not fulfilled by the current state. It is guaranteed that the transactions in the write request are sequenced adjacent to each other (with no intervention from other write requests). Only the ones with failed preconditions are left out.

For `P`, the value of a key is an object with attributes `"old"` or `"oldEmpty"` or `"isArray"`. With `"old"` one can specify a JSON value that has to be present for the condition to be fulfilled. With `"oldEmpty"`, which can take a boolean value, one can specify that the key value needs to be not set `true` or set to an arbitrary value `false`. With `"isArray"` one can specify that the value must be an array. As a shortcut, `"old"` values of scalar or array type may be stored directly in the attribute. Examples:

```
{ "/a/b/c": { "old": [1,2,3] }}
```

is a precondition specifying that the previous value of the key `"/a/b/c"` key must be `[1,2,3]`. The same could be achieved by using

```
{ "/a/b/c": [1, 2, 3] }
```

Consider the agency in initialised as above let's review the responses from the agency as follows:

```
curl -L http://server:port/write -d '[[{"/a/b/c":[1,2,3,4],"/a/b/pi"},{"/a/b/c":[1,2,3]}]]'
```

```
{
  "results": [19]
}
```

The condition is fulfilled in the first run and would be wrong in a second returning

```
{
  "results": [0]
}
```

`0` as a result means that the precondition failed and no "real" log number was returned.

```
{ "/a/e": { "oldEmpty": false } }
```

means that the value of the key `"a/e"` must be set (to something, which can be `null`!). The condition

```
{ "/a/e": { "oldEmpty": true } }
```

means that the value of the key `"a/e"` must be unset. The condition

```
{ "/a/b/c": { "isArray": true } }
```

means that the value of the key `"a/b/c"` must be an array.

The update value U is an object, the attribute names are again key strings and the values are objects with optional attributes `"new"`, `"op"` and `"ttl"`. They have the following meaning:

`"op"` determines the operation, possible values are `"set"` (the default, if left out), `"delete"`, `"increment"`, `"decrement"`, `"push"`, `"pop"`, `"shift"` or `"prepend"`

`"new"` is the new value, can be omitted for the `"delete"` operation and for `"increment"` and `"decrement"`, where `1` is implied

`"ttl"`, if present, the new value that is being set gets a time to live in seconds, given by a numeric value in this attribute. It is only guaranteed that the actual removal of the value is done according to the system clock, so up to clock skew between servers. The removal is done by an additional write transaction that is automatically generated between the regular writes.

Additional rule: If none of `"new"` and `"op"` is set or the value is not even an object, then this is to be interpreted as if it were

```
{ "op": "set", "new": <VALUE> }
```

which amounts to setting the value with no precondition.

Examples:

```
{ "/a": { "op": "set", "new": 12 } }
```

sets the value of the key `"/a"` to `12`. The same could have been achieved by

```
{ "/a": 12 }
```

or by

```
{ "/a": { "new": 12} }
```

The operation

```
{ "/a/b": { "new": { "c": [1,2,3,4] } } }
```

sets the key `"/a/b"` to `{"c": [1,2,3,4]}`. Note that in the above example this is the same as setting the value of `"/a/b/c"` to `[1,2,3,4]`. The difference is, that if `a/b` had other sub attributes, then this transaction would delete all these other attributes and make `"/a/b"` equal to `{"c": [1,2,3,4]}`, whereas setting `"/a/b/c"` to `[1,2,3,4]` would retain all attributes other than `"c"` in `"/a/b"`.

Here are some more examples for full transactions (update/precondition pairs). The transaction

```
[ { "/a/b": { "new": { "c": [1,2,3,4] } } },
  { "/a/b": { "old": { "c": [1,2,3] } } } ]
```

sets the key `"/a/b"` to `{"c":[1,2,3,4]}` if and only if it was `{"c":[1,2,3]}` before. Note that this fails if `"/a/b"` had other attributes than `"c"`. The transaction

```
[ { "/x": { "op": "delete" } },
  { "/x": { "old": false } } ]
```

clears the value of the key `"/x"` if this old value was false.

```
[ { "/y": { "new": 13 },
  { "/y": { "oldEmpty": true } } }
```

sets the value of `"/y"` to `13`, but only, if it was unset before.

```
[ { "/z": { "op": "push", "new": "Max" } } ]
```

appends the string `"Max"` to the end of the list stored in the `"z"` attribute, or creates an array `["Max"]` in `"z"` if it was unset or not an array.

```
[ { "/u": { "op": "pop" } } ]
```

removes the last entry of the array stored under `"u"`, if the value of `"u"` is not set or not an array.

HTTP-headers for write operations

`X-ArangoDB-Agency-Mode` with possible values `"waitForCommitted"`, `"waitForSequenced"` and `"noWait"`.

In the first case the write operation only returns when the commit to the replicated log has actually happened. In the second case the write operation returns when the write transactions that fulfilled their preconditions have been sequenced and thus it is known, which of the write transactions in the given array had fulfilled preconditions. In both cases the body is a JSON array containing the indexes of the transactions in the list that had fulfilled preconditions.

In the last case, `"noWait"`, the operation returns immediately, an empty body is returned. To get any information about the result of the operation one has to specify a tag (see below) and ask about the status later on.

`X-ArangoDB-Agency-Tag` with an arbitrary UTF-8 string value.

Observers

External services to the agency may announce themselves or others to be observers of arbitrary existing or future keys in the key-value-store. The agency must then inform the observing service of any changes to the subtree below the observed key. The notification is done by virtue of POST requests to a required valid URL.

In order to observe any future modification below say `"/a/b/c"`, a observer is announced through posting the below document to the agency's write REST handler:

```
[ { "/a/b/c":
    { "op":  "observe",
     "url": "http://<host>:<port>/<path>"
    }
  } ]
```

The observer is notified of any changes to that target until such time that it removes itself as an observer of that key through

```
[ { "/a/b/c":
    { "op":  "unobserve",
      "url": “http://<host>:<port>/<path>" } } ]
```

Note that the last document removes all observations from entities below `"/a/b/c"`. In particular, issuing

```
[ { "/": "unobserve", "url": "http://<host>:<port>/<path>"} ]
```

will result in the removal of all observations for URL `"http://<host>:<port>/<path>"`. The notifying POST requests are submitted immediately with any complete array of changes to the read db of the leader of create, modify and delete events accordingly; The body

```
{ "term": "5",
  "index": 167,
  "/": {
    "/a/b/c" : { "op": "modify", "old": 1, "new": 2 } },
    "/constants/euler" : {"op": "create", "new": 2.718281828459046 },
    "/constants/pi": { "op": "delete" } } }
```

HTTP Interface for Miscellaneous functions

This is an overview of ArangoDB's HTTP interface for miscellaneous functions.

Return server version

returns the server version number

`GET /_api/version`

Query Parameters

- *details* (optional): If set to *true*, the response will contain a *details* attribute with additional information about included components and their versions. The attribute names and internals of the *details* object may vary depending on platform and ArangoDB version.

Returns the server name and version number. The response is a JSON object with the following attributes:

A json document with these Properties is returned:

HTTP 200

is returned in all cases.

- **version**: the server version string. The string has the format "*major.minor.sub*". *major* and *minor* will be numeric, and *sub* may contain a number or a textual version.
- **details**: an optional JSON object with additional details. This is returned only if the *details* query parameter is set to *true* in the request.
- **server**: will always contain *arango*

Example:

Return the version information

```
shell> curl --dump - http://localhost:8529/_api/version

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "server" : "arango",
  "version" : "3.1.18",
  "license" : "community"
}
```

Example:

Return the version information with details

```
shell> curl --dump - http://localhost:8529/_api/version?details=true

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "server" : "arango",
  "version" : "3.1.18",
  "license" : "community",
  "details" : {
    "architecture" : "64bit",
    "arm" : "false",
    "asan" : "false",
    "asm-crc32" : "true",
    "assertions" : "true",
    "boost-version" : "1.62.0",
    "build-date" : "2017-04-18 17:22:09",
    "build-repository" : "heads/3.1-0-gc4d96b612d-dirty",
    "compiler" : "gcc [6.3.0 20170321]",
    "cplusplus" : "201103",
    "debug" : "false",
    "endianness" : "little",
    "failure-tests" : "false",
    "fd-client-event-handler" : "poll",
    "fd-setsize" : "1024",
    "full-version-string" : "ArangoDB 3.1.18 [linux] 64bit maintainer mode, using VPack 0.1.30, ICU 54.1, V8 5.0.71.3
9, OpenSSL 1.0.2k  26 Jan 2017",
    "icu-version" : "54.1",
    "jemalloc" : "false",
    "maintainer-mode" : "true",
    "openssl-version" : "OpenSSL 1.0.2k  26 Jan 2017",
    "platform" : "linux",
    "rocksdb-version" : "4.8.0",
    "server-version" : "3.1.18",
    "sizeof int" : "4",
    "sizeof void*" : "8",
    "sse42" : "false",
    "tcmalloc" : "false",
    "unaligned-access" : "true",
    "v8-version" : "5.0.71.39",
    "vpack-version" : "0.1.30",
    "zlib-version" : "1.2.8",
    "mode" : "server"
  }
}
```

Return Codes

- *200*: is returned in all cases.

Reply Body

- **version**: the server version string. The string has the format "*major.minor.sub*". *major* and *minor* will be numeric, and *sub* may contain a number or a textual version.
- **details**: an optional JSON object with additional details. This is returned only if the *details* query parameter is set to *true* in the request.
- **server**: will always contain *arango*

```
shell> curl --dump - http://localhost:8529/_api/version

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Return the version information with details

```
shell> curl --dump - http://localhost:8529/_api/version?details=true

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8
```

show response body

Flushes the write-ahead log

Sync the WAL to disk.

`PUT /_admin/wal/flush`

Path Parameters

- *waitForSync* (optional): Whether or not the operation should block until the not-yet synchronized data in the write-ahead log was synchronized to disk.
- *waitForCollector* (optional): Whether or not the operation should block until the data in the flushed log has been collected by the write-ahead log garbage collector. Note that setting this option to *true* might block for a long time if there are long-running transactions and the write-ahead log garbage collector cannot finish garbage collection.

Flushes the write-ahead log. By flushing the currently active write-ahead logfile, the data in it can be transferred to collection journals and datafiles. This is useful to ensure that all data for a collection is present in the collection journals and datafiles, for example, when dumping the data of a collection.

Return Codes

- *200*: Is returned if the operation succeeds.
- *405*: is returned when an invalid HTTP method is used.

Retrieves the configuration of the write-ahead log

fetch the current configuration.

`GET /_admin/wal/properties`

Retrieves the configuration of the write-ahead log. The result is a JSON object with the following attributes:

- *allowOversizeEntries*: whether or not operations that are bigger than a single logfile can be executed and stored
- *logfileSize*: the size of each write-ahead logfile
- *historicLogfiles*: the maximum number of historic logfiles to keep
- *reserveLogfiles*: the maximum number of reserve logfiles that ArangoDB allocates in the background
- *syncInterval*: the interval for automatic synchronization of not-yet synchronized write-ahead log data (in milliseconds)
- *throttleWait*: the maximum wait time that operations will wait before they get aborted if case of write-throttling (in milliseconds)
- *throttleWhenPending*: the number of unprocessed garbage-collection operations that, when reached, will activate write-throttling. A value of *0* means that write-throttling will not be triggered.

Example:

```
shell> curl --dump - http://localhost:8529/_admin/wal/properties

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "allowOversizeEntries" : true,
  "logfileSize" : 33554432,
  "historicLogfiles" : 10,
  "reserveLogfiles" : 3,
  "syncInterval" : 100,
  "throttleWait" : 15000,
  "throttleWhenPending" : 0
}
```

Return Codes

- *200*: Is returned if the operation succeeds.
- *405*: is returned when an invalid HTTP method is used.

Configures the write-ahead log

configure parameters of the wal

`PUT /_admin/wal/properties`

Configures the behavior of the write-ahead log. The body of the request must be a JSON object with the following attributes:

- *allowOversizeEntries*: whether or not operations that are bigger than a single logfile can be executed and stored
- *logfileSize*: the size of each write-ahead logfile
- *historicLogfiles*: the maximum number of historic logfiles to keep
- *reserveLogfiles*: the maximum number of reserve logfiles that ArangoDB allocates in the background
- *throttleWait*: the maximum wait time that operations will wait before they get aborted if case of write-throttling (in milliseconds)
- *throttleWhenPending*: the number of unprocessed garbage-collection operations that, when reached, will activate write-throttling. A value of *0* means that write-throttling will not be triggered.

Specifying any of the above attributes is optional. Not specified attributes will be ignored and the configuration for them will not be modified.

Example:

```
shell> curl -X PUT --data-binary @- --dump - http://localhost:8529/_admin/wal/properties <<EOF
{
  "logfileSize" : 33554432,
  "allowOversizeEntries" : true
}
EOF

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "allowOversizeEntries" : true,
  "logfileSize" : 33554432,
  "historicLogfiles" : 10,
  "reserveLogfiles" : 3,
  "syncInterval" : 100,
  "throttleWait" : 15000,
  "throttleWhenPending" : 0
}
```

Return Codes

- *200*: Is returned if the operation succeeds.
- *405*: is returned when an invalid HTTP method is used.

Returns information about the currently running transactions

returns information about the currently running transactions

`GET /_admin/wal/transactions`

Returns information about the currently running transactions. The result is a JSON object with the following attributes:

- *runningTransactions*: number of currently running transactions
- *minLastCollected*: minimum id of the last collected logfile (at the start of each running transaction). This is *null* if no transaction is running.
- *minLastSealed*: minimum id of the last sealed logfile (at the start of each running transaction). This is *null* if no transaction is running.

Example:

```
shell> curl --dump - http://localhost:8529/_admin/wal/transactions

HTTP/1.1 200 OK
content-type: application/json; charset=utf-8

{
  "runningTransactions" : 1,
  "minLastCollected" : 9456,
  "minLastSealed" : 9456
}
```

Return Codes

- *200*: Is returned if the operation succeeds.
- *405*: is returned when an invalid HTTP method is used.

Return system time

Get the current time of the system

`GET /_admin/time`

The call returns an object with the attribute *time*. This contains the current system time as a Unix timestamp with microsecond precision.

Return Codes

- *200*: Time was returned successfully.

Return current request

Send back what was sent in, headers, post body etc.

`GET /_admin/echo`

The call returns an object with the following attributes:

- *headers*: object with HTTP headers received
- *requestType*: the HTTP request method (e.g. GET)
- *parameters*: object with query parameters received

Return Codes

- *200*: Echo was returned successfully.

Return current request and continues

Send back what was sent in, headers, post body etc.

`GET /_admin/long_echo`

The call returns an object with the following attributes:

- *headers*: object with HTTP headers received
- *requestType*: the HTTP request method (e.g. GET)
- *parameters*: object with query parameters received

Return Codes

- *200*: Echo was returned successfully.

Return the required version of the database

returns the version of the database.

`GET /_admin/database/target-version`

Returns the database-version that this server requires. The version is returned in the *version* attribute of the result.

Return Codes

- *200*: Is returned in all cases.

Sleep for a specified amount of seconds

Suspend the execution for a specified duration before returnig

`GET /_admin/sleep`

Query Parameters

- *duration* (required): wait `duration` seconds until the reply is sent.

The call returns an object with the attribute *duration*. This takes as many seconds as the duration argument says.

Return Codes

- *200*: Sleep was conducted successfully.

Initiate shutdown sequence

initiates the shutdown sequence

`DELETE /_admin/shutdown`

This call initiates a clean shutdown sequence.

Return Codes

- *200*: is returned in all cases.

Runs tests on server

show the available unittests on the server.

`POST /_admin/test`

Post Body

body (required): A JSON object containing an attribute *tests* which lists the files containing the test suites.

free style json body

Executes the specified tests on the server and returns an object with the test results. The object has an attribute "error" which states whether any error occurred. The object also has an attribute "passed" which indicates which tests passed and which did not.

Execute program

Execute a script on the server.

`POST /_admin/execute`

Post Body

body (required): The body to be executed.

free style json body

Executes the javascript code in the body on the server as the body of a function with no arguments. If you have a *return* statement then the return value you produce will be returned as content type *application/json*. If the parameter *returnAsJSON* is set to *true*, the result will be a JSON object describing the return value directly, otherwise a string produced by JSON.stringify will be returned.