

Package ‘aws.s3’

June 15, 2017

Type Package

Title AWS S3 Client Package

Version 0.3.3

Date 2017-05-24

Description A simple client package for the Amazon Web Services (AWS) Simple Storage Service (S3) REST API <<https://aws.amazon.com/s3/>>.

License GPL (>= 2)

URL <https://github.com/cloudyr/aws.s3>

BugReports <https://github.com/cloudyr/aws.s3/issues>

Imports utils, tools, httr, xml2 (> 1.0.0), base64enc, digest,
aws.signature (>= 0.2.6)

Suggests testthat, datasets

RoxygenNote 6.0.1

NeedsCompilation no

Author Thomas J. Leeper [aut, cre],
Boettiger Carl [ctb],
Andrew Martin [ctb],
Mark Thompson [ctb],
Tyler Hunt [ctb],
Steven Akins [ctb],
Bao Nguyen [ctb],
Thierry Onkelinx [ctb]

Maintainer Thomas J. Leeper <thosjleeper@gmail.com>

Repository CRAN

Date/Publication 2017-06-14 22:12:27 UTC

R topics documented:

aws.s3-package	2
bucketlist	3

bucket_exists	4
copy_object	4
delete_bucket	5
delete_object	6
delete_website	7
getobject	7
get_acceleration	8
get_acl	9
get_bucket	10
get_bucketname	11
get_bucket_policy	12
get_cors	13
get_lifecycle	14
get_location	14
get_notification	15
get_object	16
get_replication	17
get_requestpayment	18
get_tagging	19
get_torrent	20
get_uploads	20
get_versions	21
put_bucket	22
put_object	23
s3HTTP	25
s3save	26
s3saveRDS	28
s3source	29
s3sync	30
s3write_using	31

Index **33**

aws.s3-package	<i>aws.s3-package</i>
----------------	-----------------------

Description

AWS S3 Client Package

Details

A simple client package for the Amazon Web Services (AWS) Simple Storage Service (S3) REST API.

Author(s)

Thomas J. Leeper <thosjleeper@gmail.com>

`bucketlist`*List Buckets*

Description

List buckets as a data frame

Usage

```
bucketlist(add_region = FALSE, ...)
```

```
bucket_list_df(add_region = FALSE, ...)
```

Arguments

`add_region` A logical (by default FALSE) indicating whether to add “Region” column to the output data frame. This simply induces a loop over [get_location](#) for each bucket.

`...` Additional arguments passed to [s3HTTP](#).

Details

`bucketlist` performs a GET operation on the base s3 endpoint and returns a list of all buckets owned by the authenticated sender of the request. If authentication is successful, this function provides a list of buckets available to the authenticated user. In this way, it can serve as a “hello world!” function, to confirm that one’s authentication credentials are working correctly.

`bucket_list_df` and `bucketlist` are identical.

Value

A data frame of buckets.

References

[API Documentation](#)

See Also

[get_bucket](#), [get_object](#)

bucket_exists	<i>Bucket exists?</i>
---------------	-----------------------

Description

Check whether a bucket exists and is accessible with the current authentication keys.

Usage

```
bucket_exists(bucket, ...)
```

Arguments

bucket	Character string with the name of the bucket, or an object of class “s3_bucket”.
...	Additional arguments passed to s3HTTP .

Value

TRUE if bucket exists and is accessible, else FALSE.

References

[API Documentation](#)

copy_object	<i>Copy Objects</i>
-------------	---------------------

Description

Copy objects between S3 buckets

Usage

```
copy_object(from_object, to_object = from_object, from_bucket, to_bucket,
  headers = list(), ...)
```

```
copy_bucket(from_bucket, to_bucket, ...)
```

Arguments

from_object	A character string containing the name the object you want to copy.
to_object	A character string containing the name the object should have in the new bucket.
from_bucket	A character string containing the name of the bucket you want to copy from.
to_bucket	A character string containing the name of the bucket you want to copy into.
headers	List of request headers for the REST call.
...	Additional arguments passed to s3HTTP .

Details

copy_object copies an object from one bucket to another without bringing it into local memory. For copy_bucket, all objects from one bucket are copied to another (limit 1000 objects). The same keys are used in the old bucket as in the new bucket.

Value

Something...

References

[API Documentation](#)

delete_bucket	<i>Delete Bucket</i>
---------------	----------------------

Description

Deletes an S3 bucket.

Usage

```
delete_bucket(bucket, ...)
```

Arguments

bucket	Character string with the name of the bucket, or an object of class "s3_bucket".
...	Additional arguments passed to s3HTTP .

Value

TRUE if successful, FALSE otherwise.

References

[API Documentation](#)

delete_object	<i>Delete object</i>
---------------	----------------------

Description

Deletes one or more objects from an S3 bucket.

Usage

```
delete_object(object, bucket, quiet = TRUE, ...)
```

Arguments

object	Character string with the object key, or an object of class “s3_object”. In most cases, if object is specified as the latter, bucket can be omitted because the bucket name will be extracted from “Bucket” slot in object.
bucket	Character string with the name of the bucket, or an object of class “s3_bucket”.
quiet	A logical indicating whether (when object is a list of multiple objects), to run the operation in “quiet” mode. Ignored otherwise. See API documentation for details.
...	Additional arguments passed to s3HTTP .

Details

object can be a single object key, an object of class “s3_object”, or a list of either.

Value

TRUE if successful, otherwise an object of class `aws_error` details if not.

References

[API Documentation](#)

See Also

[put_object](#)

delete_website	<i>Bucket Website configuration</i>
----------------	-------------------------------------

Description

Get/Put/Delete the website configuration for a bucket.

Usage

```
delete_website(bucket, ...)
```

```
put_website(bucket, request_body, ...)
```

```
get_website(bucket, ...)
```

Arguments

bucket	Character string with the name of the bucket, or an object of class “s3_bucket”.
...	Additional arguments passed to s3HTTP .
request_body	A character string containing an XML request body, as defined in the specification in the API Documentation .

Value

For `put_website` and `get_website`, a list containing the website configuration, if one has been set. For `delete_website`: TRUE if successful, FALSE otherwise. An `aws_error` object may be returned if the request failed.

References

[API Documentation: PUT website](#) [API Documentation: GET website](#) [API Documentation: DELETE website](#)

getobject	<i>Deprecated</i>
-----------	-------------------

Description

These functions are deprecated.

Usage

```

getobject(...)
saveobject(...)
headobject(...)
copyobject(...)
copybucket(...)
putbucket(...)
putobject(...)
deleteobject(...)
getbucket(...)
deletebucket(...)
bucketexists(...)

```

Arguments

... Arguments passed to updated versions of each function.

get_acceleration	<i>Bucket Acceleration</i>
------------------	----------------------------

Description

Get/Put acceleration settings or retrieve acceleration status of a bucket.

Usage

```

get_acceleration(bucket, ...)
put_acceleration(bucket, status = c("Enabled", "Suspended"), ...)

```

Arguments

bucket	Character string with the name of the bucket, or an object of class "s3_bucket".
...	Additional arguments passed to s3HTTP .
status	Character string specifying whether acceleration should be "Enabled" or "Suspended".

Details

Transfer acceleration is a AWS feature that enables potentially faster file transfers to and from S3, particularly when making cross-border transfers (such as from a European client location to the 'us-east-1' S3 region). Acceleration must be enabled before it can be used. Once enabled, `accelerate = TRUE` can be passed to any `aws.s3` function via [s3HTTP](#). `get_acceleration` returns the acceleration status of a bucket; `put_acceleration` enables or suspends acceleration.

Value

For `get_acceleration`: If acceleration has never been enabled or suspend, the value is `NULL`. Otherwise, the status is returned (either "Enabled" or "Suspended"). For `put_acceleration`: If acceleration has never been enabled or suspend, the value is `NULL`.

References

[API Documentation: PUT Bucket accelerate](#) [API Documentation: GET Bucket accelerate](#)

Examples

```
## Not run:
b <- bucketlist()
get_acceleration(b[[1]])
put_acceleration(b[[1]], "Enabled")
get_acceleration(b[[1]])
put_acceleration(b[[1]], "Suspended")

## End(Not run)
```

get_acl

Get or put bucket/object ACLs

Description

Access Control Lists (ACLs) control access to buckets and objects. These functions retrieve and modify ACLs for either objects or buckets.

Usage

```
get_acl(object, bucket, ...)

put_acl(object, bucket, body, ...)
```

Arguments

object	Character string with the object key, or an object of class "s3_object". In most cases, if object is specified as the latter, bucket can be omitted because the bucket name will be extracted from "Bucket" slot in object.
bucket	Character string with the name of the bucket, or an object of class "s3_bucket".

... Additional arguments passed to [s3HTTP](#).
 body A character string containing an XML-formatted ACL.

Details

get_acl retrieves an XML-formatted ACL for either an object (if specified) or a bucket (if specified).

Value

For get_acl a character string containing an XML-formatted ACL. For put_acl: if successful, TRUE.

get_bucket	<i>List bucket contents</i>
------------	-----------------------------

Description

List the contents of an S3 bucket as either a list or data frame

Usage

```
get_bucket(bucket, prefix = NULL, delimiter = NULL, max = NULL,
           marker = NULL, parse_response = TRUE, ...)
```

```
get_bucket_df(bucket, prefix = NULL, delimiter = NULL, max = NULL,
              marker = NULL, ...)
```

Arguments

bucket	Character string with the name of the bucket, or an object of class "s3_bucket".
prefix	Character string that limits the response to keys that begin with the specified prefix
delimiter	Character string used to group keys. Read the AWS doc for more detail.
max	Integer indicating the maximum number of keys to return. The function will recursively access the bucket in case max > 1000. Use max = Inf to retrieve all objects.
marker	Character string that specifies the key to start with when listing objects in a bucket. Amazon S3 returns object keys in alphabetical order, starting with key after the marker in order.
parse_response	logical, should we attempt to parse the response?
...	Additional arguments passed to s3HTTP .

Details

From the AWS doc: “This implementation of the GET operation returns some or all (up to 1000) of the objects in a bucket. You can use the request parameters as selection criteria to return a subset of the objects in a bucket.” The `max` and `marker` arguments can be used to retrieve additional pages of results. Values from a call are store as attributes

Value

`get_bucket` returns a list of objects in the bucket (with class “`s3_bucket`”), while `get_bucket_df` returns a data frame (the only difference is the application of the `as.data.frame()` method to the list of bucket contents. If `max` is greater than 1000, multiple API requests are executed and the attributes attached to the response object reflect only the final request.

References

[API Documentation](#)

See Also

[bucketlist](#), [get_object](#)

Examples

```
## Not run:
# basic usage
b <- bucketlist()
get_bucket(b[1,1])
get_bucket_df(b[1,1])

# bucket names with dots
## this (default) should work:
get_bucket("this.bucket.has.dots", url_style = "path")
## this probably wont:
#get_bucket("this.bucket.has.dots", url_style = "virtual")

## End(Not run)
```

get_bucketname

Utility Functions

Description

Some utility functions for working with S3 objects and buckets

Usage

```

get_bucketname(x, ...)

## S3 method for class 'character'
get_bucketname(x, ...)

## S3 method for class 's3_bucket'
get_bucketname(x, ...)

## S3 method for class 's3_object'
get_bucketname(x, ...)

```

Arguments

x	An object.
...	Ignored.

Value

get_bucketname returns a character string.

get_bucket_policy	<i>Bucket policies</i>
-------------------	------------------------

Description

Get/Put/Delete the bucket access policy for a bucket.

Usage

```

get_bucket_policy(bucket, parse_response = TRUE, ...)

put_bucket_policy(bucket, policy, ...)

delete_bucket_policy(bucket, ...)

```

Arguments

bucket	Character string with the name of the bucket, or an object of class “s3_bucket”.
parse_response	A logical indicating whether to return the response as is, or parse and return as a list. Default is FALSE.
...	Additional arguments passed to s3HTTP .
policy	A character string containing a bucket policy.

Details

Bucket policies regulate who has what access to a bucket and its contents. The header argument can be used to specify “canned” policies and [put_bucket_policy](#) can be used to specify a more complex policy. The [AWS Policy Generator](#) can be useful for creating the appropriate JSON policy structure.

Value

For `get_policy`: A character string containing the JSON representation of the policy, if one has been set. For `delete_policy` and `put_policy`: TRUE if successful, FALSE otherwise.

References

[API Documentation](#) [API Documentation](#) [AWS Policy Generator](#)

get_cors	<i>CORS</i>
----------	-------------

Description

Get/Put/Delete the cross origin resource sharing configuration information for a bucket.

Usage

```
get_cors(bucket, ...)
```

```
put_cors(bucket, ...)
```

```
delete_cors(bucket, ...)
```

Arguments

bucket Character string with the name of the bucket, or an object of class “s3_bucket”.

... Additional arguments passed to [s3HTTP](#).

Value

For `get_cors`: A list with cors configuration and rules. For `delete_cors`: TRUE if successful, FALSE otherwise.

References

[API Documentation: PUT cors](#) [API Documentation: GET cords](#) [API Documentation: DELETE cors](#)

get_lifecycle	<i>Lifecycle</i>
---------------	------------------

Description

Get/Put/Delete the lifecycle configuration information for a bucket.

Usage

```
get_lifecycle(bucket, ...)
```

```
put_lifecycle(bucket, request_body, ...)
```

```
delete_lifecycle(bucket, ...)
```

Arguments

bucket	Character string with the name of the bucket, or an object of class “s3_bucket”.
...	Additional arguments passed to s3HTTP .
request_body	A character string containing an XML request body, as defined in the specification in the API Documentation .

Value

For `get_lifecycle`: a list with lifecycle configuration, if it has been configured. For `delete_lifecycle`: TRUE if successful, FALSE otherwise.

References

[API Documentation: PUT lifecycle API Documentation: GET lifecycle API Documentation: DELETE lifecycle](#)

get_location	<i>Bucket location</i>
--------------	------------------------

Description

Get the AWS region location of bucket.

Usage

```
get_location(bucket, ...)
```

Arguments

bucket Character string with the name of the bucket, or an object of class “s3_bucket”.
... Additional arguments passed to [s3HTTP](#).

Value

A character string containing the region, if one has been set.

References

[API Documentation](#)

get_notification *Notifications*

Description

Get/put the notification configuration for a bucket.

Usage

```
get_notification(bucket, ...)  
put_notification(bucket, request_body, ...)
```

Arguments

bucket Character string with the name of the bucket, or an object of class “s3_bucket”.
... Additional arguments passed to [s3HTTP](#).
request_body A character string containing an XML request body, as defined in the specification in the [API Documentation](#).

Value

A list containing the notification configuration, if one has been set.

References

[API Documentation: GET](#) [API Documentation: PUT](#)

get_object

Get object

Description

Retrieve an object from an S3 bucket

Usage

```
get_object(object, bucket, headers = list(), parse_response = FALSE, ...)
```

```
save_object(object, bucket, file, headers = list(), ...)
```

```
head_object(object, bucket, ...)
```

Arguments

object	Character string with the object key, or an object of class “s3_object”. In most cases, if object is specified as the latter, bucket can be omitted because the bucket name will be extracted from “Bucket” slot in object.
bucket	Character string with the name of the bucket, or an object of class “s3_bucket”.
headers	List of request headers for the REST call.
parse_response	Passed through to s3HTTP , as this function requires a non-default setting. There is probably no reason to ever change this.
...	Additional arguments passed to s3HTTP .
file	An R connection, or file name specifying the local file to save the object into.

Details

`get_object` retrieves an object into memory as a raw vector. `save_object` saves an object to a local file. `head_object` checks whether an object exists by executing an HTTP HEAD request; this can be useful for checking object headers such as “content-length” or “content-type”.

Some users may find the raw vector response format of `get_object` unfamiliar. The object will also carry attributes, including “content-type”, which may be useful for deciding how to subsequently process the vector. Two common strategies are as follows. For text content types, running [charToRaw](#) may be the most useful first step to make the response human-readable. Alternatively, converting the raw vector into a connection using [rawConnection](#) may also be useful, as that can often then be passed to parsing functions just like a file connection would be.

Value

If `file = NULL`, a raw object. Otherwise, a character string containing the file name that the object is saved to.

References

[API Documentation: GET Object](#)

[API Documentation: GET Object torrent](#)

[API Document: HEAD Object](#)

See Also

[get_bucket](#), [put_object](#), [delete_object](#)

Examples

```
## Not run:
# get an object in memory
## create bucket
b <- put_bucket("myexamplebucket")

## save a dataset to the bucket
s3save(mtcars, bucket = b, object = "mtcars")
obj <- get_bucket(b)
## get the object in memory
x <- get_object(obj[[1]])
load(rawConnection(x))
"mtcars" %in% ls()

# save an object locally
y <- save_object(obj[[1]], file = object[[1]][["Key"]])
y %in% dir()

# return object using 'S3 URI' syntax
get_object("s3://myexamplebucket/mtcars")

# return parts of an object
## use 'Range' header to specify bytes
get_object(object = obj[[1]], headers = list('Range' = 'bytes=1-120'))

## End(Not run)
```

get_replication

Bucket replication

Description

Get/Delete the replication configuration for a bucket.

Usage

```
get_replication(bucket, ...)
```

```
put_replication(bucket, request_body, ...)
```

```
delete_replication(bucket, ...)
```

Arguments

bucket	Character string with the name of the bucket, or an object of class “s3_bucket”.
...	Additional arguments passed to s3HTTP .
request_body	A character string containing an XML request body, as defined in the specification in the API Documentation .

Details

`get_replication` gets the current replication policy. `delete_replication` deletes the replication policy for a bucket.

Value

For `get_replication`: A list containing the replication configuration, if one has been set. For `delete_replication`: TRUE if successful, FALSE otherwise.

References

[API Documentation: PUT replication](#) [API Documentation: GET replication](#) [API Documentation: DELETE replication](#)

<code>get_requestpayment</code>	<i>requestPayment</i>
---------------------------------	-----------------------

Description

Get/Put the `requestPayment` subresource for a bucket.

Usage

```
get_requestpayment(bucket, ...)
```

```
put_requestpayment(bucket, ...)
```

Arguments

bucket	Character string with the name of the bucket, or an object of class “s3_bucket”.
...	Additional arguments passed to s3HTTP .

Value

A list containing the requestPayment information, if set.

References

[API Documentation](#)

get_tagging	<i>Bucket tagging</i>
-------------	-----------------------

Description

Get/delete the tag set for a bucket.

Usage

```
get_tagging(bucket, ...)

put_tagging(bucket, tags = list(), ...)

delete_tagging(bucket, ...)
```

Arguments

bucket	Character string with the name of the bucket, or an object of class “s3_bucket”.
...	Additional arguments passed to s3HTTP .
tags	A list containing key-value pairs of tag names and values.

Value

A list containing the tag set, if one has been set. For delete_tagging: TRUE if successful, FALSE otherwise.

References

[API Documentation: PUT tagging](#) [API Documentation: GET tagging](#) [API Documentation: DELETE tagging](#)

Examples

```
## Not run:
put_tagging("mybucket", tags = list(foo = "1", bar = "2"))
get_tagging("mybucket")
delete_tagging("mybucket")

## End(Not run)
```

get_torrent	<i>Get object torrent</i>
-------------	---------------------------

Description

Retrieves a Bencoded dictionary (BitTorrent) for an object from an S3 bucket.

Usage

```
get_torrent(object, bucket, ...)
```

Arguments

object	Character string with the object key, or an object of class “s3_object”. In most cases, if object is specified as the latter, bucket can be omitted because the bucket name will be extracted from “Bucket” slot in object.
bucket	Character string with the name of the bucket, or an object of class “s3_bucket”.
...	Additional arguments passed to s3HTTP .

Value

Something.

References

[API Documentation](#)

get_uploads	<i>Multipart uploads</i>
-------------	--------------------------

Description

Get a list of multipart uploads for a bucket.

Usage

```
get_uploads(bucket, ...)
```

Arguments

bucket	Character string with the name of the bucket, or an object of class “s3_bucket”.
...	Additional arguments passed to s3HTTP .

Value

A list containing the multipart upload information.

References[API Documentation](#)

get_versions	<i>Bucket versions</i>
--------------	------------------------

Description

Get/Put versioning settings or retrieve versions of bucket objects.

Usage

```
get_versions(bucket, ...)
```

```
get_versioning(bucket, ...)
```

```
put_versioning(bucket, status = c("Enabled", "Suspended"), ...)
```

Arguments

bucket	Character string with the name of the bucket, or an object of class "s3_bucket".
...	Additional arguments passed to s3HTTP .
status	Character string specifying whether versioning should be "Enabled" or "Suspended".

Details

get_versioning returns the versioning status of a bucket; put_versioning sets the versioning status. get_versions returns information about bucket versions.

Value

For get_versioning: If versioning has never been enabled or suspend, the value is NULL. Otherwise, the status is returned (either "Enabled" or "Suspended"). For put_versioning: If versioning has never been enabled or suspend, the value is NULL. Otherwise, the status is returned (either "Enabled" or "Suspended"). For get_versions: A list.

References[API Documentation](#) [API Documentation](#) [API Documentation](#)

Examples

```
## Not run:
put_versioning("mybucket")
get_versioning("mybucket")
get_versions("mybucket")

## End(Not run)
```

put_bucket

Create bucket

Description

Creates a new S3 bucket.

Usage

```
put_bucket(bucket, region = Sys.getenv("AWS_DEFAULT_REGION"),
  acl = c("private", "public-read", "public-read-write", "aws-exec-read",
    "authenticated-read", "bucket-owner-read", "bucket-owner-full-control"),
  headers = list(), ...)
```

Arguments

bucket	Character string with the name of the bucket, or an object of class “s3_bucket”.
region	A character string containing the AWS region. If missing, defaults to value of environment variable AWS_DEFAULT_REGION.
acl	A character string indicating a “ canned ” access control list . By default all bucket contents and objects therein are given the ACL “private”. This can later be viewed using get_acl and modified using put_acl .
headers	List of request headers for the REST call.
...	Additional arguments passed to s3HTTP .

Details

Bucket policies regulate who has what access to a bucket and its contents. The header argument can be used to specify “canned” policies and [put_bucket_policy](#) can be used to specify a more complex policy. The [AWS Policy Generator](#) can be useful for creating the appropriate JSON policy structure.

Value

TRUE if successful.

References

[API Documentation AWS Policy Generator](#)

See Also

[bucketlist](#), [get_bucket](#), [delete_bucket](#), [put_object](#)

Examples

```
## Not run:
  put_bucket("examplebucket")

# set a "canned" ACL to, e.g., make bucket publicly readable
  put_bucket("examplebucket", headers = list(`x-amz-acl` = "public-read"))

## End(Not run)
```

put_object	<i>Put object</i>
------------	-------------------

Description

Puts an object into an S3 bucket

Usage

```
put_object(file, object, bucket, multipart = FALSE, acl = c("private",
  "public-read", "public-read-write", "aws-exec-read", "authenticated-read",
  "bucket-owner-read", "bucket-owner-full-control"), headers = list(), ...)

put_folder(folder, bucket, ...)
```

Arguments

file	A character string containing the filename (or full path) of the file you want to upload to S3. Alternatively, a raw vector containing the file can be passed directly, in which case object needs to be specified explicitly.
object	A character string containing the name the object should have in S3 (i.e., its "object key"). If missing, the filename is used.
bucket	Character string with the name of the bucket, or an object of class "s3_bucket".
multipart	A logical indicating whether to use multipart uploads. See http://docs.aws.amazon.com/AmazonS3/latest/dev/mpuoverview.html . If file is less than 100 MB, this is ignored.
acl	A character string indicating a "canned" access control list. By default all bucket contents and objects therein are given the ACL "private". This can later be viewed using get_acl and modified using put_acl .

headers	List of request headers for the REST call.
...	Additional arguments passed to s3HTTP .
folder	A character string containing a folder name. (A trailing slash is not required.)

Details

This provide a generic interface for sending files (or serialized, in-memory representations thereof) to S3. Some convenience wrappers are provided for common tasks: e.g., [s3save](#) and [s3saveRDS](#).

Note that S3 is a flat file store. So there is no folder hierarchy as in a traditional hard drive. However, S3 allows users to create pseudo-folders by prepending object keys with `foldername/`. The `put_folder` function is provided as a high-level convenience function for creating folders. This is not actually necessary as objects with slashes in their key will be displayed in the S3 web console as if they were in folders, but it may be useful for creating an empty directory (which is possible in the web console).

Value

If successful, TRUE.

References

[API Documentation](#)

See Also

[put_bucket](#), [get_object](#), [delete_object](#)

Examples

```
## Not run:
library("datasets")

# write file to S3
tmp <- tempfile()
on.exit(unlink(tmp))
utils::write.csv(mtcars, file = tmp)
put_object(tmp, object = "mtcars.csv", bucket = "myexamplebucket")

# create a "folder" in a bucket
put_folder("example", bucket = "myexamplebucket")
## write object to the "folder"
put_object(tmp, object = "example/mtcars.csv", bucket = "myexamplebucket")

# write serialized, in-memory object to S3
x <- rawConnection(raw(0), "w")
utils::write.csv(mtcars, x)
put_object(rawConnectionValue(x), object = "mtcars.csv", bucket = "myexamplebucketname")

# use `headers` for server-side encryption
## require appropriate bucket policy
```



```

put_object(file = tmp, object = "mtcars.csv", bucket = "myexamplebucket",
           headers = c('x-amz-server-side-encryption' = 'AES256'))

# alternative "S3 URI" syntax:
put_object(rawConnectionValue(x), object = "s3://myexamplebucketname/mtcars.csv")
close(x)

# read the object back from S3
read.csv(text = rawToChar(get_object(object = "s3://myexamplebucketname/mtcars.csv")))

## End(Not run)

```

s3HTTP

S3 HTTP Requests

Description

This is the workhorse function for executing API requests for S3.

Usage

```

s3HTTP(verb = "GET", bucket = "", path = "", query = NULL,
       headers = list(), request_body = "", accelerate = FALSE,
       dualstack = FALSE, parse_response = TRUE, check_region = TRUE,
       url_style = c("path", "virtual"), base_url = "s3.amazonaws.com",
       verbose = getOption("verbose", FALSE),
       region = Sys.getenv("AWS_DEFAULT_REGION", "us-east-1"),
       key = Sys.getenv("AWS_ACCESS_KEY_ID"),
       secret = Sys.getenv("AWS_SECRET_ACCESS_KEY"),
       session_token = Sys.getenv("AWS_SESSION_TOKEN"), ...)

```

Arguments

verb	A character string containing an HTTP verb, defaulting to “GET”.
bucket	A character string with the name of the bucket, or an object of class “s3_bucket”. If the latter and a region can be inferred from the bucket object attributes, then that region is used instead of region.
path	A character string with the name of the object to put in the bucket (sometimes called the object or ‘key name’ in the AWS documentation.)
query	Any query arguments, passed as a named list of key-value pairs.
headers	A list of request headers for the REST call.
request_body	A character string containing request body data.
accelerate	A logical indicating whether to use AWS transfer acceleration, which can produce significant speed improvements for cross-country transfers. Acceleration only works with buckets that do not have dots in bucket name.

dualstack	A logical indicating whether to use “dual stack” requests, which can resolve to either IPv4 or IPv6. See http://docs.aws.amazon.com/AmazonS3/latest/dev/dual-stack-endpoints.html .
parse_response	A logical indicating whether to return the response as is, or parse and return as a list. Default is TRUE.
check_region	A logical indicating whether to check the value of region against the apparent bucket region. This is useful for avoiding (often confusing) out-of-region errors. Default is TRUE.
url_style	A character string specifying either “path” (the default), or “virtual”-style S3 URLs.
base_url	A character string specifying the base URL for the request. There is no need to set this, as it is provided only to generalize the package to (potentially) support S3-compatible storage on non-AWS servers.
verbose	A logical indicating whether to be verbose. Default is given by options(“verbose”).
region	A character string containing the AWS region. Ignored if region can be inferred from bucket. If missing, defaults to “us-east-1”.
key	A character string containing an AWS Access Key ID. If missing, defaults to value stored in environment variable “AWS_ACCESS_KEY_ID”.
secret	A character string containing an AWS Secret Access Key. If missing, defaults to value stored in environment variable “AWS_SECRET_ACCESS_KEY”.
session_token	Optionally, a character string containing an AWS temporary Session Token. If missing, defaults to value stored in environment variable “AWS_SESSION_TOKEN”.
...	Additional arguments passed to an HTTP request function. such as GET .

Details

This is mostly an internal function for executing API requests. In almost all cases, users do not need to access this directly.

Value

the S3 response, or the relevant error.

s3save

save/load

Description

Save/load R object(s) to/from S3

Usage

```
s3save(..., object, bucket, envir = parent.frame(), opts = NULL)
```

```
s3save_image(object, bucket, opts = NULL)
```

```
s3load(object, bucket, envir = parent.frame(), ...)
```

Arguments

...	For s3save, one or more R objects to be saved via save and uploaded to S3. For s3load, see opts.
object	For s3save, a character string of the name of the object you want to save to. For s3load, a character string of the name of the object you want to load from S3.
bucket	Character string with the name of the bucket, or an object of class "s3_bucket".
envir	For s3save, an R environment to save objects from; for s3load, the environment to load objects into. Default is the parent.frame() from which the function is called.
opts	Additional arguments passed to s3HTTP .

Value

For s3save, a logical, invisibly. For s3load, NULL invisibly.

References

[API Documentation](#)

See Also

[s3saveRDS](#), [s3readRDS](#)

Examples

```
## Not run:
# create bucket
b <- put_bucket("myexamplebucket")

# save a dataset to the bucket
s3save(mtcars, iris, object = "somedata.Rdata", bucket = b)
get_bucket(b)

# load the data from bucket
e <- new.env()
s3load(object = "somedata.Rdata", bucket = b, envir = e)
ls(e)

# cleanup
rm(e)
delete_object(object = "somedata.Rdata", bucket = "myexamplebucket")
delete_bucket("myexamplebucket")

## End(Not run)
```

s3saveRDS

*saveRDS/readRDS***Description**

Serialization interface to read/write R objects to S3

Usage

```
s3saveRDS(x, object = paste0(as.character(substitute(x)), ".rds"), bucket,
          compress = TRUE, ...)
```

```
s3readRDS(object, bucket, ...)
```

Arguments

x	For s3saveRDS, a single R object to be saved via saveRDS and uploaded to S3. x is analogous to the object argument in saveRDS .
object	Character string with the object key, or an object of class "s3_object". In most cases, if object is specified as the latter, bucket can be omitted because the bucket name will be extracted from "Bucket" slot in object.
bucket	Character string with the name of the bucket, or an object of class "s3_bucket".
compress	A logical. See saveRDS .
...	Additional arguments passed to s3HTTP .

Details

Note that early versions of s3saveRDS from `aws.s3 <= 0.2.4` unintentionally serialized objects to big endian format (due to defaults in [serialize](#)). This can create problems when attempting to read these files using [readRDS](#). The function attempts to catch the issue and read accordingly, but may fail. The solution used internally is `unserialize(memDecompress(get_object(), "gzip"))`

Value

For s3saveRDS, a logical. For s3readRDS, an R object.

Author(s)

Steven Akins <skawesome@gmail.com>

See Also

[s3save](#), [s3load](#)

Examples

```
## Not run:
# create bucket
b <- put_bucket("myexamplebucket")

# save a single object to s3
s3saveRDS(x = mtcars, bucket = "myexamplebucket", object = "mtcars.rds")

# restore it under a different name
mtcars2 <- s3readRDS(object = "mtcars.rds", bucket = "myexamplebucket")
identical(mtcars, mtcars2)

# cleanup
delete_object(object = "mtcars.rds", bucket = "myexamplebucket")
delete_bucket("myexamplebucket")

## End(Not run)
```

s3source

Source from S3

Description

Source R code (a la [source](#)) from S3

Usage

```
s3source(object, bucket, ..., opts = NULL)
```

Arguments

object	Character string with the object key, or an object of class “s3_object”. In most cases, if object is specified as the latter, bucket can be omitted because the bucket name will be extracted from “Bucket” slot in object.
bucket	Character string with the name of the bucket, or an object of class “s3_bucket”.
...	Additional arguments passed to s3HTTP .
opts	Additional arguments passed to get_object for retrieving the R syntax file.

Value

See [source](#)

See Also

[s3saveRDS](#), [s3save](#), [get_object](#)

Examples

```
## Not run:
# create bucket
b <- put_bucket("myexamplebucket")

# save some code to the bucket
cat("x <- 'hello world!'\n", file = "example.R")
put_object("example.R", object = "example.R", bucket = b)
get_bucket(b)

# source the code from the bucket
s3source(object = "example.R", bucket = b, echo = TRUE)

# cleanup
unlink("example.R")
delete_object(object = "example.R", bucket = b)
delete_bucket("myexamplebucket")

## End(Not run)
```

s3sync

S3 file sync

Description

Sync files/directories to/from S3

Usage

```
s3sync(files = dir(recursive = TRUE), bucket, verbose = TRUE, ...)
```

Arguments

files	A character vector specifying relative paths to files to be synchronized. The default is all files in the working directory and subdirectories.
bucket	Character string with the name of the bucket, or an object of class "s3_bucket".
verbose	A logical indicating whether to be verbose (the default is TRUE).
...	Additional arguments passed to s3HTTP .

Details

s3sync synchronizes specified files to an S3 bucket. This works best if a local directory (and its subdirectories) correspond directly to the contents of an S3 bucket. If the bucket does not exist, it is created. Similarly, if local directories do not exist (corresponding to leading portions of object keys), they are created, recursively. Object keys are generated based on files and local files are named (and organized into directories) based on object keys. A slash is interpreted as a directory level. Local objects are copied to S3 and S3 objects are copied locally. This copying is performed

conditionally. Objects existing locally but not in S3 are uploaded using `put_object`. Objects existing in S3 but not locally, are saved using `save_object`. If objects exist in both places, the MD5 checksum for each is compared; when identical, no copying is performed. If the checksums differ, local files are replaced with the bucket version if the local file is older and the S3 object is replaced if the local file is newer. If checksums differ but modified times match (which seems unlikely), a warning is issued.

Value

A logical.

References

[aws s3 sync command line](#)

See Also

[get_bucket](#), [put_object](#), [save_object](#)

Examples

```
## Not run:
s3sync(bucket = "examplebucket")

## End(Not run)
```

s3write_using

Custom read and write

Description

Read/write objects from/to S3 using a custom function

Usage

```
s3write_using(x, FUN, ..., object, bucket, opts = NULL)
```

```
s3read_using(FUN, ..., object, bucket, opts = NULL)
```

Arguments

x	For <code>s3write_using</code> , a single R object to be saved via the first argument to <code>FUN</code> and uploaded to S3.
FUN	For <code>s3write_using</code> , a function to which <code>x</code> and a file path will be passed (in that order).
...	Additional arguments to <code>FUN</code>

object	Character string with the object key, or an object of class “s3_object”. In most cases, if object is specified as the latter, bucket can be omitted because the bucket name will be extracted from “Bucket” slot in object.
bucket	Character string with the name of the bucket, or an object of class “s3_bucket”.
opts	Optional additional arguments passed to put_object or save_object , respectively.

Value

For `s3write_using`, a logical, invisibly. For `s3read_using`, the output of FUN applied to the file from object.

See Also

[s3saveRDS](#), [s3readRDS](#), [put_object](#), [get_object](#)

Examples

```
## Not run:
library("datasets")
# create bucket
b <- put_bucket("myexamplebucket")

# save a dataset to the bucket as a csv
if (require("utils")) {
  s3write_using(mtcars, FUN = write.csv, object = "mtcars.csv", bucket = b)
}

# load dataset from the bucket as a csv
if (require("utils")) {
  s3read_using(FUN = read.csv, object = "mtcars.csv", bucket = b)
}

# cleanup
delete_object(object = "mtcars.csv", bucket = b)
delete_bucket(bucket = b)

## End(Not run)
```


Index

- *Topic **package**
 - aws.s3-package, 2
- *Topic **service**
 - bucketlist, 3
- aws.s3 (aws.s3-package), 2
- aws.s3-package, 2
- bucket_exists, 4
- bucket_list_df (bucketlist), 3
- bucketexists (getobject), 7
- bucketlist, 3, 11, 23
- charToRaw, 16
- copy_bucket (copy_object), 4
- copy_object, 4
- copybucket (getobject), 7
- copyobject (getobject), 7
- delete_bucket, 5, 23
- delete_bucket_policy
 - (get_bucket_policy), 12
- delete_cors (get_cors), 13
- delete_lifecycle (get_lifecycle), 14
- delete_object, 6, 17, 24
- delete_replication (get_replication), 17
- delete_tagging (get_tagging), 19
- delete_website, 7
- deletebucket (getobject), 7
- deleteobject (getobject), 7
- GET, 26
- get_acceleration, 8
- get_acl, 9, 22, 23
- get_bucket, 3, 10, 17, 23, 31
- get_bucket_df (get_bucket), 10
- get_bucket_policy, 12
- get_bucketname, 11
- get_cors, 13
- get_lifecycle, 14
- get_location, 3, 14
- get_notification, 15
- get_object, 3, 11, 16, 24, 29, 32
- get_replication, 17
- get_requestpayment, 18
- get_tagging, 19
- get_torrent, 20
- get_uploads, 20
- get_versioning (get_versions), 21
- get_versions, 21
- get_website (delete_website), 7
- getbucket (getobject), 7
- getobject, 7
- head_object (get_object), 16
- headobject (getobject), 7
- put_acceleration (get_acceleration), 8
- put_acl, 22, 23
- put_acl (get_acl), 9
- put_bucket, 22, 24
- put_bucket_policy, 13, 22
- put_bucket_policy (get_bucket_policy), 12
- put_cors (get_cors), 13
- put_folder (put_object), 23
- put_lifecycle (get_lifecycle), 14
- put_notification (get_notification), 15
- put_object, 6, 17, 23, 23, 31, 32
- put_replication (get_replication), 17
- put_requestpayment
 - (get_requestpayment), 18
- put_tagging (get_tagging), 19
- put_versioning (get_versions), 21
- put_website (delete_website), 7
- putbucket (getobject), 7
- putobject (getobject), 7
- rawConnection, 16
- readRDS, 28
- s3HTTP, 3–10, 12–16, 18–22, 24, 25, 27–30

s3load, [28](#)
s3load (s3save), [26](#)
s3read_using (s3write_using), [31](#)
s3readRDS, [27](#), [32](#)
s3readRDS (s3saveRDS), [28](#)
s3save, [24](#), [26](#), [28](#), [29](#)
s3save_image (s3save), [26](#)
s3saveRDS, [24](#), [27](#), [28](#), [29](#), [32](#)
s3source, [29](#)
s3sync, [30](#)
s3write_using, [31](#)
save, [27](#)
save_object, [31](#), [32](#)
save_object (get_object), [16](#)
saveobject (getobject), [7](#)
saveRDS, [28](#)
serialize, [28](#)
source, [29](#)