Package 'AnVIL'

March 29, 2021

Title Bioconductor on the AnVIL compute environment

Version 1.2.0

Description The AnVIL is cloud computing resource developed in part
by the National Human Genome Research Institute. The AnVIL package
provides end-user and devloper functionality. For the end-user,
AnVIL provides fast binary package installation, utitlities for
working with Terra / AnVIL table and data resources, and
convenient functions for file movement to and from Google cloud
storage. For developers, AnVIL provides programatic access to the
Terra, Leonardo, Dockstore, and Gen3 RESTful programming
interface, including helper functions to transform JSON responses
to more formats more amenable to manipulation in R.

License Artistic-2.0

Encoding UTF-8

Depends R (>= 3.6), dplyr

Imports stats, utils, methods, futile.logger, jsonlite, httr, curl,
rapiclient (>= 0.1.3), tibble, tidyselect, rlang, BiocManager

Suggests knitr, testthat, withr, readr

Collate utilities.R authenticate.R Service.R Services.R Leonardo.R
Terra.R Dockstore.R Gen3.R Response.R gcloud_sdk.R gcloud.R
gsutil.R localize.R av.R avworkflow.R drs.R zzz.R

LazyData true

VignetteBuilder knitr

biocViews Infrastructure

RoxygenNote 7.1.1

git_url https://git.bioconductor.org/packages/AnVIL

git_branch RELEASE_3_12

git_last_commit c686619

git_last_commit_date 2020-10-27

Date/Publication 2021-03-29

Author Martin Morgan [aut, cre],
Nitesh Turaga [aut],
BJ Stubbs [ctb],
Vincent Carey [ctb],

Marcel Ramos [ctb],
Sweta Gopaulakrishnan [ctb],
Valerie Obenchain [ctb]

Maintainer Martin Morgan <mtmorgan.bioc@gmail.com>

R topics documented:

---

avworkflow_jobs   *Working with AnVIL Workflow submissions and file outputs*

Description

avworkflow_jobs() returns a tibble summarizing submitted workflow jobs for a namespace and name.

avworkflow_files() returns a tibble containing information and file paths to workflow outputs.

Usage

```
avworkflow_jobs(namespace = avworkspace_namespace(), name = avworkspace_name())

avworkflow_files(submissionId = NULL, bucket = avbucket())
```

Arguments

namespace   character(1) AnVIL workspace namespace as returned by, e.g., avworkspace_namespace()

name   character(1) AnVIL workspace name as returned by, eg., avworkspace_name().

submissionId   a character() of workflow submission ids, or a tibble with column submissionId, or NULL / missing. See 'Details'.

bucket   character(1) name of the google bucket in which the workflow products are available, as gs://.... Usually the bucket of the active workspace, returned by avbucket().

Details

For avworkflow_files(), the submissionId is the identifier associated with the workflow job, and is present in the return value of avworkflow_jobs(); the example illustrates how the first row of avworkflow_jobs() (i.e., the most recenltly completed workflow) can be used as input to avworkflow_files(). When submissionId is not provided, the return value is for the most recently submitted workflow of the namespace and name of avworkspace().

Value

avworkflow_jobs() returns a tibble, sorted by submissionDate, with columns

- submissionId character() job identifier from the workflow runner.
- submitter character() AnVIL user id of individual submitting the job.
- submissionDate POSIXct() date (in local time zone) of job submission.
- succeeded integer() number of workflows succeeding.
- failed integer() number of workflows failing.

avworkflow_files() returns a tibble with columns

- file: character() 'base name' of the file in the bucket.
- workflow: character() name of the workflow the file is associated with.
- task: character() name of the task in the workflow that generated the file.
- path: charcter() full path to the file in the google bucket.

Examples

```
if (gcloud_exists() && nzchar(avworkspace_name()))
    ## from within AnVIL
    avworkflow_jobs()

if (gcloud_exists() && nzchar(avworkspace_name())) {
    ## e.g., from within AnVIL
    avworkflow_jobs() %>%
    ## select most recent workflow
    head(1) %>%
    ## find paths to output and log files on the bucket
    avworkflow_files()
}
```

---

avworkspace_namespace   *Functions for convenient user interaction with AnVIL resources*

Description

avworkspace_namespace() and avworkspace_name() are utiliity functions to retrieve workspace namespace and name from environment variables or interfaces usually available in AnVIL notebooks or RStudio sessions. avworkspace() provides a convenient way to specify workspace namespace and name in a single command.

'avtables()' describes tables available in a workspace. Tables can be visualized under the DATA tab, TABLES item. 'avtable()' returns an AnVIL table. 'avtable_import()' imports a data.frame to an AnVIL table. 'avtable_import_set()' imports set membership (i.e., a subset of an existing table) information to an AnVIL table. 'avtable_delete_values()' removes rows from an AnVIL table.

'avdata()' returns key-value tables representing the information visualized under the DATA tab, 'REFERENCE DATA' and 'OTHER DATA' items.

'avbucket()' returns the workspace bucket, i.e., the google bucket associated with a workspace. Bucket content can be visualized under the 'DATA' tab, 'Files' item.

'avfiles_ls()' returns the paths of files in the workspace bucket. 'avfiles_backup()' copies files from the compute node file system to the workspace bucket. 'avfiles_restore()' copies files from the workspace bucket to the compute node file system. 'avfiles_rm()' removes files or directories from the workspace bucket.

avruntimes() returns a tibble containing information about runtimes (notebooks or RStudio instances, for example) that the current user has access to.

'avdisks()' returns a tibble containing information about persistent disks associatd with the current user.

Usage

```
avworkspace_namespace(namespace = NULL)

avworkspace_name(name = NULL)

avworkspace(workspace = NULL)

avtables(namespace = avworkspace_namespace(), name = avworkspace_name())

avtable(table, namespace = avworkspace_namespace(), name = avworkspace_name())

avtable_import(
  .data,
  entity = names(.data)[[1]],
  namespace = avworkspace_namespace(),
  name = avworkspace_name()
)

avtable_import_set(
  .data,
  origin,
  set = names(.data)[[1]],
  member = names(.data)[[2]],
  namespace = avworkspace_namespace(),
  name = avworkspace_name()
)

avtable_delete_values(
  table,
  values,
  namespace = avworkspace_namespace(),
  name = avworkspace_name()
)

avdata(namespace = avworkspace_namespace(), name = avworkspace_name())

avbucket(
  namespace = avworkspace_namespace(),
  name = avworkspace_name(),
  as_path = TRUE
)
```

```
avfiles_ls(
  path = "",
  full_names = FALSE,
  recursive = FALSE,
  namespace = avworkspace_namespace(),
  name = avworkspace_name()
)

avfiles_backup(
  source,
  destination = "",
  recursive = FALSE,
  parallel = TRUE,
  namespace = avworkspace_namespace(),
  name = avworkspace_name()
)

avfiles_restore(
  source,
  destination = ".",
  recursive = FALSE,
  parallel = TRUE,
  namespace = avworkspace_namespace(),
  name = avworkspace_name()
)

avfiles_rm(
  source,
  recursive = FALSE,
  parallel = TRUE,
  namespace = avworkspace_namespace(),
  name = avworkspace_name()
)

avruntimes()

avdisks()
```

Arguments

namespace	character(1) AnVIL workspace namespace as returned by, e.g., `avworkspace_namespace()`
name	character(1) AnVIL workspace name as returned by, eg., `avworkspace_name()`.
workspace	when present, a 'character(1)' providing the concatenated namespace and name, e.g., '"bioconductor-rpci-anvil/Bioconductor-Package-AnVIL"'
table	character(1) table name as returned by, e.g., 'avtables()'.
.data	A tibble or data.frame for import as an AnVIL table.
entity	'character(1)' column name of '.data' to be used as imported table name. When the table comes from R, this is usually a column name such as 'sample'. The data will be imported into AnVIL as a table 'sample', with the 'sample' column included with suffix '_id', e.g., 'sample_id'. A column in '.data' with suffix '_id' can also be used, e.g., 'entity = "sample_id"', creating the table 'sample'

	with column 'sample_id' in AnVIL. Finally, a value of 'entity' that is not a column in '.data', e.g., 'entity = "unknown"', will cause a new table with name 'entity' and entity values 'seq_len(nrow(.data))'.
origin	character(1) name of the entity (table) used to create the set e.g "sample", "participant", etc.
set	'character(1)' column name of '.data' identifying the set(s) to be created.
member	'character()' vector of entity from the avtable identified by 'origin'. The values may repeat if an ID is in more than one set
values	vector of values in the entity (key) column of 'table' to be deleted. A table 'sample' has an associated entity column with suffix '_id', e.g., 'sample_id'. Rows with entity column entries matching 'values' are deleted.
as_path	logical(1) when TRUE (default) return bucket with prefix 'gs://' (for 'avbucket()') or 'gs://<bucket-id>' (for 'avfiles_ls()').
path	For 'avfiles_ls(), the character(1) file or directory path to list.
	For 'avfiles_rm()', the character() (perhaps with length greater than 1) of files or directory paths to be removed. The elements of 'path' can contain glob-style patterns, e.g., 'vign*'.
full_names	logical(1) return names relative to 'path' ('FALSE', default) or root of the workspace bucket?
recursive	logical(1) list files recursively?
source	character() file paths. for 'avfiles_backup()', 'source' can include directory names when 'recursive = TRUE'.
destination	character(1) a google bucket ('gs://<bucket-id>/...') to write files. The default is the workspace bucket.
parallel	logical(1) backup files using parallel transfer? See '?gsutil_cp()'.

Details

avworkspace_namespace() is the billing account. If the namespace= argument is not provided, try gcloud_project(), and if that fails try Sys.getenv("WORKSPACE_NAMESPACE").

```
`avworkspace_name()` is the name of the workspace as it appears
in \url{https://app.terra.bio/#workspaces}. If not provided,
`avworkspace_name()` tries to use
`Sys.getenv("WORKSPACE_NAME")`.
```

```
Values are cached across sessions, so explicitly providing
`avworkspace_*()` is required at most once per session. Revert
to system settings with arguments `NA`.
```

'avtable_import_set()' creates new rows in a table '<origin>_set'. One row will be created for each distinct value in the column identified by 'set'. Each row entry has a corresponding column '<origin>' linking to one or more rows in the '<origin>' table, as given in the 'member' column. The operation is somewhat like 'split(member, set)'.

'avfiles_backup()' can be used to back-up individual files or entire directories, recursively. When 'recursive = FALSE', files are backed up to the bucket with names approximately 'paste0(destination, "/", basename(source))'. When 'recursive = TRUE' and source is a directory 'path/to/foo/', files are backed up to bucket names that include the directory name, approximately 'paste0(destination, "/",

dir(basename(source), full.names = TRUE))'. Naming conventions are described in detail in 'gsutil_help("cp")'.

'avfiles_restore()' behaves in a manner analogous to 'avfiles_backup()', copying files from the workspace bucket to the compute node file system.

Value

avworkspace_namespace(), and avworkspace_name() return character(1) identifiers.

'avworkspace()' returns the character(1) concatenated namespace and name.

'avtables()': A tibble with columns identifying the table, the number of records, and the column names.

'avtable()': a tibble of data corresponding to the AnVIL table 'table' in the specified workspace.

'avtable_import()' returns a 'character(1)' name of the imported AnVIL tibble.

'avtable_import_set()' returns a 'character(1)' name of the imported AnVIL tibble.

'avtable_delete_values()' returns a 'tibble' representing deleted entities, invisibly.

'avdata()' returns a tibble with five columns: '"type"' represents the origin of the data from the 'REFERENCE' or 'OTHER' data menus. '"table"' is the table name in the 'REFERENCE' menu, or 'workspace' for the table in the 'OTHER' menu, the key used to access the data element, the value label associated with the data element and the value (e.g., google bucket) of the element.

'avbucket()' returns a 'character(1)' bucket identifier, prefixed with 'gs://' if 'as_path = TRUE'.

'avfiles_ls()' returns a character vector of files in the workspace bucket.

'avfiles_backup()' returns, invisibly, the status code of the 'gsutil_cp()' command used to back up the files.

'avfiles_rm()' on success, returns a list of the return codes of 'gsutil_rm()', invisibly.

avruntimes() returns a tibble with columns

- id: integer() runtime identifier.
- googleProject: character() billing account.
- tool: character() e.g., "Jupyter", "RStudio".
- status character() e.g., "Stopped", "Running".
- creator character() AnVIL account, typically "user@gmail.com".
- createdDate character() creation date.
- destroyedDate character() destruction date, or NA.
- dateAccessed character() date of (first?) access.
- runtimeName character().
- clusterServiceAccount character() service ('pet') account for this runtime.
- masterMachineType character() It is unclear which 'tool' populates which of the machineType columns).
- workerMachineType character().
- machineType character().
- persistentDiskId integer() identifier of persistent disk (see avdisks()), or NA.

avdisks() returns a tibble with columns

- id character() disk identifier.

- googleProject: character() billing account.

- status, e.g, "Ready"

- size integer() in GB.

- diskType character().

- blockSize integer().

- creator character() AnVIL account, typically "user@gmail.com".

- createdDate character() creation date.

- destroyedDate character() destruction date, or NA.

- dateAccessed character() date of (first?) access.

- zone character() e.g.. "us-central1-a".

- name character().

Examples

```
avworkspace_namespace()

avworkspace_name()

avworkspace()

## Not run:
## editable copy of '1000G-high-coverage-2019' workspace
avworkspace("bioconductor-rpci-anvil/1000G-high-coverage-2019")
sample <-
    avtable("sample") %>%                                # existing table
    mutate(set = sample(head(LETTERS), nrow(.), TRUE))  # arbitrary groups
sample %>%                                      # new 'participant_set' table
    avtable_import_set("participant", "set", "participant")
sample %>%                                      # new 'sample_set' table
    avtable_import_set("sample", "set", "name")

## End(Not run)

if (gcloud_exists() && nzchar(avworkspace_name()))
    ## from within AnVIL
    avdata()

if (gcloud_exists() && nzchar(avworkspace_name()))
    ## From within AnVIL...
    bucket <- avbucket()                       # discover bucket

## Not run:
path <- file.path(bucket, "mtcars.tab")
gsutil_ls(dirname(path))                   # no 'mtcars.tab'...
write.table(mtcars, gsutil_pipe(path, "w")) # write to bucket
gsutil_stat(path)                          # yep, there!
read.table(gsutil_pipe(path, "r"))         # read from bucket

## End(Not run)
if (gcloud_exists() && nzchar(avworkspace_name()))
    avfiles_ls()

## Not run:
```

```
## backup all files in the current directory
## default buckets are gs://<bucket-id>/<file-names>
avfiles_backup(dir())
## backup working directory, recursively
## default buckets are gs://<bucket-id>/<basename(getwd())>/...
avfiles_backup(getwd(), recursive = TRUE)

## End(Not run)

if (gcloud_exists())
    ## from within AnVIL
    avruntimes()

if (gcloud_exists())
    ## from within AnVIL
    avdisks()
```

---

drs_stat   *Manipulate 'DRS' (Data Repository Service) URLs*

Description

drs_stat() resolves zero or more DRS URLs to their google bucket location.

drs_cp() copies 0 or more DRS URIs to a google bucket or local folder

Usage

```
drs_stat(source = character())

drs_cp(source, destination, ...)
```

Arguments

source	character() DRS URLs (beginning with 'drs://') to resources managed by the 'martha' DRS resolution server.
destination	character(1) directory path in which to retrieve files.
...	additional arguments, passed to gsutil_cp() for file copying.

Details

drs_stat() uses the AnVIL 'pet' account associated with a runtime. The pet account is discovered by default when evaluated on an AnVIL runtime (e.g., in RStudio or a Jupyter notebook in the AnVIL), or can be found in the return value of avruntimes().

Value

drs_stat() returns a tbl with the following columns:

- fileName: character() (resolver sometimes returns null).
- size: integer() (resolver sometimes returns null).
- contentType: character() (resolver sometimes returns null).

- gsUri: character() (resolver sometimes returns null).

- timeCreated: character() (the time created formatted using ISO 8601; resolver sometimes returns null).

- timeUpdated: character() (the time updated formatted using ISO 8601; resolver sometimes returns null).

- bucket: character() (resolver sometimes returns null).

- name: character() (resolver sometimes returns null).

- googleServiceAccount: list() (null unless the DOS url belongs to a Bond supported host).

- hashes: list() (contains the hashes type and their checksum value; if unknown. it returns null)

drs_cp() returns a tibble like drs_stat(), but with additional columns

- simple: logical() value indicating whether resolution used a simple signed URL (TRUE) or auxilliary service account.

- destination: character() full path to retrieved object(s)

Examples

```
drs_eg_hca <- paste0(
    "drs://drs.data.humancellatlas.org/",
    "4cf48dbf-cf09-452e-bb5b-fd016af0c747?version=2019-09-14T024754.281908Z"
)

drs_eg_anvil <- c(
    "drs://dg.ANV0/975bd45f-f022-4fad-b9a2-3a00c3b8792c",
    "drs://dg.ANV0/00008531-03d7-418c-b3d3-b7b22b5381a0"
)

if (gcloud_exists())
    # no pet account needed for HCA data
    drs_stat(drs_eg_hca)

if (gcloud_exists() && startsWith(gcloud_account(), "pet-")) {
    ## from within AnVIL
    drs_stat(drs_eg_anvil)
}

if (gcloud_exists()) {
    destination <- tempfile()
    dir.create(destination)
    tbl <- drs_cp(drs_eg_hca, destination)
    readLines(tbl$destination, warn = FALSE)
}
```

---

gcloud   *Interact with the gcloud command line utility*

Description

These functions invoke the 'gcloud' command line utility. See gsutil for details on how 'gcloud' is located.

'gcloud_exists()' tests whether the 'gcloud()' command can be found on this system. See 'Details' section of 'gsutil' for where the application is searched.

'gcloud_account()': report the current gcloud account via 'gcloud config get-value account'.

'gcloud_project()': report the current gcloud project via 'gcloud config get-value project'.

'gcloud_help()': queries 'gcloud' for help for a command or sub-comand via 'gcloud help ...'.

'gcloud_cmd()' allows arbitrary 'gcloud' command execution via 'gcloud ...'. Use pre-defined functions in preference to this.

Usage

```
gcloud_exists()

gcloud_account(account = NULL)

gcloud_project(project = NULL)

gcloud_help(...)

gcloud_cmd(cmd, ...)
```

Arguments

account	character(1) Google account (e.g., 'user@gmail.com') to use for authentication.
project	character(1) billing project name.
...	Additional arguments appended to gcloud commands.
cmd	'character(1)' representing a command used to evaluate 'gcloud cmd ...'.

Value

'gcloud_exists()' returns 'TRUE' when the 'gcloud' application can be found, FALSE otherwise.

'gcloud_account()' returns a 'character(1)' vector containing the active gcloud account, typically a gmail email address.

'gcloud_project()' returns a 'character(1)' vector containing the active gcloud project.

'gcloud_help()' returns an unquoted 'character()' vector representing the text of the help manual page returned by 'gcloud help ...'.

'gcloud_cmd()' returns a 'character()' vector representing the text of the output of 'gcloud cmd ...'

Examples

```
gcloud_exists()

if (gcloud_exists())
    gcloud_account()

if (gcloud_exists())
    gcloud_help()
```

---

gsutil   *Interact with the gsutil command line utility*

Description

These functions invoke the 'gsutil' command line utility. See the "Details:" section if you have gsutil installed but the package cannot find it.

'gsutil_requesterpays()': does the google bucket require that the requester pay for access?

'gsutil_ls()': List contents of a google cloud bucket or, if 'source' is missing, all Cloud Storage buckets under your default project ID

'gsutil_exists()': check if the bucket or object exists.

'gsutil_stat()': print, as a side effect, the status of a bucket, directory, or file.

'gsutil_cp()': copy contents of 'source' to 'destination'. At least one of 'source' or 'destination' must be Google cloud bucket; 'source' can be a character vector with length greater than 1. Use 'gsutil_help("cp")' for 'gsutil' help.

'gsutil_rm()': remove contents of a google cloud bucket.

'gsutil_rsync()': synchronize a source and a destination.

'gsutil_cat()': concatenate bucket objects to standard output

'gsutil_help()': print 'man' page for the 'gsutil' command or subcommand. Note that only commandes documented on this R help page are supported.

'gsutil_pipe()': create a pipe to read from or write to a gooogle bucket object.

Usage

```
gsutil_requesterpays(source)

gsutil_ls(source = character(), ..., recursive = FALSE)

gsutil_exists(source)

gsutil_stat(source)

gsutil_cp(source, destination, ..., recursive = FALSE, parallel = TRUE)

gsutil_rm(source, ..., force = FALSE, recursive = FALSE, parallel = TRUE)

gsutil_rsync(
  source,
  destination,
  ...,
  dry = TRUE,
  delete = FALSE,
  recursive = FALSE,
  parallel = TRUE
)

gsutil_cat(source, ..., header = FALSE, range = integer())
```

```
gsutil_help(cmd = character(0))

gsutil_pipe(source, open = "r", ...)
```

Arguments

source	'character(1)', ('character()' for 'gsutil_requesterpays()', 'gsutil_ls()', 'gsutil_exists()', 'gsutil_cp()') paths to a google storage bucket, possibly with wild-cards for file-level pattern matching.
...	additional arguments passed as-is to the 'gsutil' subcommand.
recursive	'logical(1)'; perform operation recursively from 'source'?. Default: 'FALSE'.
destination	'character(1)', google cloud bucket or local file system destination path.
parallel	'logical(1)', perform parallel multi-threaded / multi-processing (default is 'TRUE').
force	'logical(1)': continue silently despite errors when removing multiple objects. Default: 'FALSE'.
dry	'logical(1)', when 'TRUE' (default), return the consequences of the operation without actually performing the operation.
delete	'logical(1)', when 'TRUE', remove files in 'destination' that are not in 'source'. Exercise caution when you use this option: it's possible to delete large amounts of data accidentally if, for example, you erroneously reverse source and destination.
header	'logical(1)' when 'TRUE' annotate each
range	(optional) 'integer(2)' vector used to form a range from-to of bytes to concatenate. 'NA' values signify concatenation from the start (first position) or to the end (second position) of the file.
cmd	'character()' (optional) command name, e.g., '"ls"' for help.
open	'character(1)' either '"r"' (read) or '"w"' (write) from the bucket.

Details

The 'gsutil' system command is required. The search for 'gsutil' starts with environment variable 'GCLOUD_SDK_PATH' providing a path to a directory containing a 'bin' directory containingin 'gsutil', 'gcloud', etc. The path variable is searched for first as an 'option()' and then system variable. If no option or global variable is found, 'Sys.which()' is tried. If that fails, 'gsutil' is searched for on defined paths. On Windows, the search tries to find 'Google\Cloud SDK\google-cloud-sdk\bin\gsutil.cmd' in the 'LOCAL APP DATA', 'Program Files', and 'Program Files (x86)' directories. On linux / macOS, the search continues with '~/google-cloud-sdk'.

'gsutil_rsync()': To make '"gs://mybucket/data"' match the contents of the local directory '"data"' you could do:

```
gsutil_rsync("data","gs://mybucket/data",delete = TRUE)
```

To make the local directory "data" the same as the contents of gs://mybucket/data:

```
gsutil_rsync("gs://mybucket/data","data",delete = TRUE)
```

If 'destination' is a local path and does not exist, it will be created.

Value

'gsutil_requesterpays()': named 'logical()' vector TRUE when requester-pays is enabled.

'gsutil_ls()': 'character()' listing of 'source' content.

'gsutil_exists()': logical(1) TRUE if bucket or object exists.

'gsutil_stat()': 'character()' description of status of objects matching 'source'.

'gsutil_cp()': exit status of 'gsutil_cp()', invisibly.

'gsutil_rm()': exit status of 'gsutil_rm()', invisibly.

'gsutil_rsync()': exit status of 'gsutil_rsync()', invisbly.

'gsutil_cat()' returns the content as a character vector.

'gsutil_help()': 'character()' help text for subcommand 'cmd'.

'gsutil_pipe()' an unopened R 'pipe()'; the mode is *not* specified, and the pipe must be used in the appropriate context (e.g., a pipe created with 'open = "r"' for input as 'read.csv()')

Examples

```
    src <- "gs://genomics-public-data/1000-genomes/other/sample_info/sample_info.csv"
if (gcloud_exists())
    gsutil_requesterpays(src) # FALSE -- no cost download

if (gcloud_exists()) {
    gsutil_exists(src)
    gsutil_stat(src)
    gsutil_ls(dirname(src))
}

if (gcloud_exists())
   gsutil_cp(src, tempdir())

if (gcloud_exists())
    gsutil_help("ls")

if (gcloud_exists()) {
    df <- read.csv(gsutil_pipe(src), 5L)
    class(df)
    dim(df)
    head(df)
}
```

---

localize   *Copy packages, folders, or files to or from google buckets.*

Description

'localize()': recursively synchronizes files from a Google storage bucket ('source') to the local file system ('destination'). This command acts recursively on the 'source' directory, and does not delete files in 'destination' that are not in 'source.

'delocalize()': synchronize files from a local file system ('source') to a Google storage bucket ('destination'). This command acts recursively on the 'source' directory, and does not delete files in 'destination' that are not in 'source'.

'install()': install R / Bioconductor packages, using fast pre-built 'binary' libraries if available.

'repositories()': repositories to search for binary (if available), Bioconductor, and CRAN packages.

'add_libpaths()': Add local library paths to '.libPaths()'.

Usage

```
localize(source, destination, dry = TRUE)

delocalize(source, destination, unlink = FALSE, dry = TRUE)

install(
  pkgs,
  lib = .libPaths()[1],
  ...,
  version = BiocManager::version(),
  binary_base_url = "https://storage.googleapis.com",
  verbose = getOption("verbose")
)

repositories(
  version = BiocManager::version(),
  binary_base_url = "https://storage.googleapis.com"
)

add_libpaths(paths)
```

Arguments

source	'character(1)', a google storage bucket or local file system directory location.
destination	'character(1)', a google storage bucket or local file system directory location.
dry	'logical(1)', when 'TRUE' (default), return the consequences of the operation without actually performing the operation.
unlink	'logical(1)' remove (unlink) the file or directory in 'source'. Default: 'FALSE'.
pkgs	'character()' packages to install from binary repository.
lib	'character(1)' library path (directory) in which to install 'pkgs'; defaults to '.libPaths()[1]'.
...	additional arguments, passed to 'install.packages()'.
version	'character(1)' or 'package_version' Bioconductor version, e.g., "3.12".
binary_base_url	
	'character(1)' host and base path for binary package 'CRAN-style' repository; not usually required by the end-user.
verbose	'logical(1)' report on package installation progress?
paths	'character()': vector of directories to add to '.libPaths()'. Paths that do not exist will be created.

Details

'repositories()' prepends an additional repository URI to 'BiocManager::repositories()'. The URI is formed by concatenating 'binary_base_url', the environment variables 'TERRA_R_PLATFORM'

and the 'major' and 'minor' components of 'TERRA_R_PLATFORM_BINARY_VERSION' and 'BiocManager::version()'. The URI is only prepended if a CRAN-style repostiory exists at that location, with binary package tar.gz content described by 'src/contrib/PACKAGES.gz'.

Value

'localize()': exit status of function 'gsutil_rsync()'.

'delocalize()': exit status of function 'gsutil_rsync()'

'install()': return value of 'install.packages()'.

'repositories()': character() of binary (if available), Bioconductor, and CRAN repositories.

'add_libpaths()': updated .libPaths(), invisibly.

Examples

```
## Not run: install(c('BiocParallel', 'BiocGenerics'))

repositories()

## Not run: add_libpaths("/tmp/host-site-library")
```

Response   *Process service responses to tibble and other data structures.*

Description

Process service responses to tibble and other data structures.

Usage

```
flatten(x)

## S4 method for signature 'response'
str(object)

## S3 method for class 'response'
as.list(x, ..., as = c("text", "raw", "parsed"))
```

Arguments

x	A 'response' object returned by the service.
object	A 'response' object returned by the service.
...	not currently used
as	character(1); one of 'raw', 'text', 'parsed'

Value

'flatten()' returns a 'tibble' where each row correseponds to a top-level list element of the return value, and columns are the unlisted second and more-nested elements.

'str()' displays a compact representation of the list-like JSON response; it returns 'NULL'.

'as.list()' retruns the content of the web service request as a list.

Examples

```
if (gcloud_exists()) {
    leonardo <- Leonardo()
    leonardo$listClusters() %>% flatten()
}

if (gcloud_exists())
   leonardo$getSystemStatus() %>% str()

if (gcloud_exists())
    leonardo$getSystemStatus() %>% as.list()
```

Service	*RESTful service constructor*

Description

RESTful service constructor

Usage

```
Service(
  service,
  host,
  config = httr::config(),
  authenticate = TRUE,
  api_url = character(),
  package = "AnVIL",
  schemes = "https",
  api_reference_url = api_url,
  api_reference_md5sum = character()
)
```

Arguments

service	character(1) The 'Service' class name, e.g., '"terra"'.
host	character(1) host name that provides the API resource, e.g., '"api.firecloud.org"'.
config	httr::config() curl options
authenticate	logical(1) use credentials from authentication service file 'auth.json' in the specified package?
api_url	optional character(1) url location of OpenAPI '.json' or '.yaml' service definition.
package	character(1) (default 'AnVIL') The package where 'api.json' yaml and (optionally) 'auth.json' files are located.
schemes	character(1) (default 'https') Specifies the transfer protocol supported by the API service.
api_reference_url	
	character(1) path to reference API. See Details.
api_reference_md5sum	
	character(1) the result of 'tools::md5sum()' applied to the reference API.

Details

This function creates a RESTful interface to a service provided by a host, e.g., "api.firecloud.org". The function requires an OpenAPI '.json' or '.yaml' specifcation as well as an (optional) '.json' authentication token. These files are located in the source directory of a pacakge, at '<pack-age>/inst/service/<service>/api.json' and '<package>/inst/service/<service>/auth.json', or at 'api_url'.

When provided, the 'api_reference_md5sum' is used to check that the file described at 'api_reference_url' has the same checksum as an author-validated version.

The service is usually a singleton, created at the package level during '.onLoad()'.

Value

An object of class `Service`.

Examples

```
.MyService <- setClass("MyService", contains = "Service")

MyService <- function() {
    .MyService(Service("my_service", host="my.api.org"))
}
```

Services	*RESTful services useful for AnVIL developers*

Description

RESTful services useful for AnVIL developers

Usage

```
empty_object

operations(x, ..., .deprecated = FALSE)

schemas(x)

tags(x, .tags, .deprecated = FALSE)

## S4 method for signature 'Service'
x$name

Leonardo()

Terra()

Dockstore()

Gen3Fence()

Gen3Indexd()
```

```
Gen3Sheepdog()

Gen3Peregrine()
```

Arguments

x	A 'Service' instance, usually a singleton provided by the package and documented on this page, e.g., 'leonardo' or 'terra'.
...	additional arguments passed to methods or, for 'operations,Service-method', to the internal 'get_operation()' function.
.deprecated	optional logical(1) include deprecated operations?
.tags	optional character() of tags to use to filter operations.
name	A symbol representing a defined operation, e.g., 'leonardo$listClusters()'.

Value

'empty_object' returns a representation to be used as arguments in function calls expecting the empty json object '{}'.

'Leonardo()' creates the API of the Leonard container deployment service at https://notebooks.firecloud.org/api-docs.yaml.

'Terra()' creates the API of the Terra cloud computational environemnt at https://api.firecloud.org/.

'Dockstore()' represents the API of the Dockstore platform to share Docker-based tools in CWL or WDL or Nextflow at https://dockstore.org

'gen3_*' APIs are not fully implemented, because a service endpoint has not been identifiied.

'Gen3Fence()' returns the authentication API at https://raw.githubusercontent.com/uc-cdis/fence/master/openapis/swagg

'Gen3Indexd()' returns the indexing service API documented at https://raw.githubusercontent.com/uc-cdis/indexd/master/openapis/swagger.yaml

'Gen3Sheepdog' returns the submission services API at https://raw.githubusercontent.com/uc-cdis/sheepdog/master/ope

'Gen3Peregrine' returns the graphQL query services API at https://raw.githubusercontent.com/uc-cdis/peregrine/master/openapis/swagger.yaml

Examples

```
empty_object

if (gcloud_exists())
    Leonardo()

if (gcloud_exists()) {
    terra <- Terra()
    tags(terra)
    tags(terra, "Billing")
}

Dockstore()
```

Index