

Package: Rarr (via r-universe)

May 29, 2026

Title Read Zarr Files in R

Version 2.0.1

Description The Zarr specification defines a format for chunked, compressed, N-dimensional arrays. Its design allows efficient access to subsets of the stored array, and supports both local and cloud storage systems. Rarr aims to implement this specification in R with minimal reliance on an external tools or libraries.

License MIT + file LICENSE

URL <https://huber-group-embl.github.io/Rarr/>,
<https://github.com/Huber-group-EMBL/Rarr>

BugReports <https://github.com/Huber-group-EMBL/Rarr/issues>

Depends R (>= 4.1.0)

Imports curl, jsonlite, lifecycle, paws.storage, R.utils, utils

Suggests BiocStyle, knitr, rmarkdown, testthat (>= 3.0.0), withr, ZarrArray

VignetteBuilder knitr

biocViews DataImport

Encoding UTF-8

Roxygen list(markdown = TRUE)

RoxygenNote 7.3.3

SystemRequirements GNU make

Config/testthat/edition 3

Config/pak/sysreqs make libxml2-dev libssl-dev

Repository <https://bioc-release.r-universe.dev>

Date/Publication 2026-05-29 09:25:29 UTC

RemoteUrl <https://github.com/bioc/Rarr>

RemoteRef RELEASE_3_23

RemoteSha 4ed1f2e9abc0393fc588911914b31b512c3f800a

Contents

compressors	2
create_empty_zarr_array	3
read_zarr_array	5
read_zarr_attributes	6
update_zarr_array	7
write_zarr_array	8
write_zarr_attributes	9
zarr_overview	10
ZarrArray	11
Index	13

compressors	<i>Define compression tool and settings</i>
-------------	---

Description

These functions select a compression tool and its setting when writing a Zarr file

Usage

```
use_blosc(cname = "lz4")
```

```
use_zlib(level = 6L)
```

```
use_gzip(level = 6L)
```

```
use_bz2(level = 6L)
```

```
use_lzma(level = 9L)
```

```
use_lz4()
```

```
use_zstd(level = 3)
```

Arguments

cname Blosc is a 'meta-compressor' providing access to several compression algorithms. This argument defines which compression tool should be used. Valid options are: 'lz4', 'lz4hc', 'blosclz', 'zstd', 'zlib', 'snappy'.

level Specify the compression level to use. The range of possible values is dependant on the compression tool being used. For example, for `use_zlib()` this argument can be between 1 & 9, while for `use_zstd()` the valid range is 1 to 22.

Value

A list containing the details of the selected compression tool. This will be written to the .zarray metadata when the Zarr array is created.

Examples

```
## define 2 compression filters for blosc (using snappy) and bzip2 (level 5)
blosc_with_snappy_compression <- use_blosc(cname = "snappy")
bzip2_compression <- use_bz2(level = 5)

## create an example array to write to a file
x <- array(runif(n = 1000, min = -10, max = 10), dim = c(10, 20, 5))

## write the array to two files using each compression filter
blosc_path <- tempfile()
bzip2_path <- tempfile()
write_zarr_array(
  x = x, zarr_array_path = blosc_path, chunk_dim = c(2, 5, 1),
  compressor = blosc_with_snappy_compression
)
write_zarr_array(
  x = x, zarr_array_path = bzip2_path, chunk_dim = c(2, 5, 1),
  compressor = bzip2_compression
)

## the contents of the two arrays should be the same
identical(read_zarr_array(blosc_path), read_zarr_array(bzip2_path))

## the size of the files on disk are not the same
sum(file.size(list.files(blosc_path, full.names = TRUE)))
sum(file.size(list.files(bzip2_path, full.names = TRUE)))
```

```
create_empty_zarr_array
```

Create an (empty) Zarr array

Description

Create an (empty) Zarr array

Usage

```
create_empty_zarr_array(
  zarr_array_path,
  dim,
  chunk_dim,
  data_type,
  order = "F",
```

```

compressor = use_zstd(),
fill_value,
nchar = NULL,
dimension_separator = if (zarr_version == 2) "." else "/",
dimension_names = NULL,
zarr_version = 3
)

```

Arguments

zarr_array_path	Character vector of length 1 giving the path to the new Zarr array.
dim	Dimensions of the new array. Should be a numeric vector with the same length as the number of dimensions.
chunk_dim	Dimensions of the array chunks. Should be a numeric vector with the same length as the dim argument.
data_type	Character vector giving the data type of the new array. Valid options are: "integer", "double", "character", "logical", which are based on standard R data types. You can also use the analogous Numpy formats: "i1", "<i2", "<i4", "<f4", "<f8", "iS", "l1". If this argument isn't provided the fill_value will be used to determine the datatype.
order	Define the layout of the bytes within each chunk. Valid options are 'column', 'row', 'F' & 'C'. 'column' or 'F' will specify "column-major" ordering, which is how R arrays are arranged in memory. 'row' or 'C' will specify "row-major" order.
compressor	What (if any) compression tool should be applied to the array chunks. The default is to use zstd compression. Supplying NULL will disable chunk compression. See compressors for more details.
fill_value	The default value for uninitialized portions of the array. Does not have to be provided, in which case the default for the specified data type will be used.
nchar	For datatype = "character" this parameter gives the maximum length of the stored strings. It is an error not to specify this for a character array, but it is ignored for other data types.
dimension_separator	The character used to separate the dimensions in the names of the chunk files. Valid options are limited to "." and "/".
dimension_names	Optional character vector with the same length as dim.
zarr_version	The version of the Zarr specification to use. Currently, either 2 or 3. The default is 3.

Value

If successful returns (invisibly) TRUE. However this function is primarily called for the side effect of initialising a Zarr array location and creating the .zarray metadata.

See Also

[write_zarr_array\(\)](#), [update_zarr_array\(\)](#)

Examples

```
new_zarr_array <- file.path(tempdir(), "temp.zarr")
create_empty_zarr_array(new_zarr_array,
  dim = c(10, 20), chunk_dim = c(2, 5),
  data_type = "integer"
)
```

read_zarr_array	<i>Read a Zarr array</i>
-----------------	--------------------------

Description

Read a Zarr array

Usage

```
read_zarr_array(zarr_array_path, index, s3_client)
```

Arguments

zarr_array_path	Path to a Zarr array. A character vector of length 1. This can either be a location on a local file system or the URI to an array in S3 storage.
index	A list of the same length as the number of dimensions in the Zarr array. Each entry in the list provides the indices in that dimension that should be read from the array. Setting a list entry to NULL will read everything in the associated dimension. If this argument is missing the entirety of the the Zarr array will be read.
s3_client	Object created by paws.storage::s3() . Only required for a file on S3. Leave as NULL for a file on local storage.

Value

An array with the same number of dimensions as the input array. The extent of each dimension will correspond to the length of the values provided to the `index` argument.

Examples

```
## Using a local file provided with the package
## This array has 3 dimensions
z1 <- system.file("extdata", "zarr_examples", "row-first", "int32.zarr", package = "Rarr")

## read the entire array
read_zarr_array(zarr_array_path = z1)

## extract values for first 10 rows, all columns, first slice
read_zarr_array(zarr_array_path = z1, index = list(1:10, NULL, 1))

## using a Zarr file hosted on Amazon S3
## This array has a single dimension with length 2729077
z2 <- "https://noaa-nwm-retro-v2-zarr-pds.s3.amazonaws.com/feature_id/"

## read the entire array
read_zarr_array(zarr_array_path = z2)

## read alternating elements
read_zarr_array(zarr_array_path = z2, index = list(seq(1, 576, 2)))
```

read_zarr_attributes *Read the attributes associated with a Zarr array or group*

Description

Read the attributes associated with a Zarr array or group

Usage

```
read_zarr_attributes(
  zarr_path,
  s3_client,
  missing = c("ignore", "warning", "error")
)
```

Arguments

zarr_path	A character vector of length 1. This provides the path to a Zarr array or group of arrays. This can either be on a local file system or on S3 storage.
s3_client	A list representing an S3 client. This should be produced by <code>paws.storage::s3()</code> .
missing	A character vector of length 1. This determines the behaviour when no file containing attributes is found. This can be one of: <ul style="list-style-type: none"> • "ignore" (the default): an empty list is returned silently • "warning": a warning is issued and an empty list is returned. • "error": an error is raised.

Value

A list containing the attributes. If the file containing attributes (.zattrs for Zarr v2 or zarr.json for Zarr v3) exists but no attributes are provided, an empty list is returned.

update_zarr_array	<i>Update (a subset of) an existing Zarr array</i>
-------------------	--

Description

Update (a subset of) an existing Zarr array

Usage

```
update_zarr_array(zarr_array_path, x, index)
```

Arguments

zarr_array_path	Character vector of length 1 giving the path to the Zarr array that is to be modified.
x	The R array (or object that can be coerced to an array) that will be written to the Zarr array.
index	A list with the same length as the number of dimensions of the target array. This argument indicates which elements in the target array should be updated.

Value

The function is primarily called for the side effect of writing to disk. Returns (invisibly) TRUE if the array is successfully updated.

Examples

```
## first create a new, empty, Zarr array
new_zarray_array <- file.path(tempdir(), "new_array.zarr")
create_empty_zarr_array(
  zarr_array_path = new_zarray_array, dim = c(20, 10),
  chunk_dim = c(10, 5), data_type = "double"
)

## create a matrix smaller than our Zarr array
small_matrix <- matrix(runif(6), nrow = 3)

## insert the matrix into the first 3 rows, 2 columns of the Zarr array
update_zarr_array(new_zarray_array, x = small_matrix, index = list(1:3, 1:2))

## reading back a slightly larger subset,
## we can see only the top left corner has been changed
read_zarr_array(new_zarray_array, index = list(1:5, 1:5))
```

<code>write_zarr_array</code>	<i>Write an R array to Zarr</i>
-------------------------------	---------------------------------

Description

Write an R array to Zarr

Usage

```
write_zarr_array(
  x,
  zarr_array_path,
  chunk_dim,
  data_type = storage.mode(x),
  order = "F",
  compressor = use_zstd(),
  fill_value,
  nchar,
  dimension_separator = if (zarr_version == 2) "." else "/",
  zarr_version = 3
)
```

Arguments

<code>x</code>	The R array that will be written to the Zarr array.
<code>zarr_array_path</code>	Character vector of length 1 giving the path to the new Zarr array.
<code>chunk_dim</code>	Dimensions of the array chunks. Should be a numeric vector with the same length as the <code>dim</code> argument.
<code>data_type</code>	Character vector giving the data type of the new array. Valid options are: "integer", "double", "character", "logical", which are based on standard R data types. You can also use the analogous Numpy formats: "i1", "<i2", "<i4", "<f4", "<f8", "iS", "l1". If this argument isn't provided the <code>fill_value</code> will be used to determine the datatype.
<code>order</code>	Define the layout of the bytes within each chunk. Valid options are 'column', 'row', 'F' & 'C'. 'column' or 'F' will specify "column-major" ordering, which is how R arrays are arranged in memory. 'row' or 'C' will specify "row-major" order.
<code>compressor</code>	What (if any) compression tool should be applied to the array chunks. The default is to use <code>zstd</code> compression. Supplying <code>NULL</code> will disable chunk compression. See compressors for more details.
<code>fill_value</code>	The default value for uninitialized portions of the array. Does not have to be provided, in which case the default for the specified data type will be used.

nchar	For character arrays this parameter gives the maximum length of the stored strings. If this argument is not specified the array provided to x will be checked and the length of the longest string found will be used so no data are truncated. However this may be slow and providing a value to nchar can provide a modest performance improvement.
dimension_separator	The character used to separate the dimensions in the names of the chunk files. Valid options are limited to "." and "/".
zarr_version	The version of the Zarr specification to use. Currently, either 2 or 3. The default is 3.

Value

The function is primarily called for the side effect of writing to disk. Returns (invisibly) TRUE if the array is successfully written.

Note

If x has dimnames, names(dimnames(x)) will be stored as the dimension_names field in the Zarr metadata.

Examples

```
new_zarr_array <- file.path(tempdir(), "integer.zarr")
x <- array(1:50, dim = c(10, 5))
write_zarr_array(
  x = x, zarr_array_path = new_zarr_array,
  chunk_dim = c(2, 5)
)
```

write_zarr_attributes *Read the .zattrs file associated with a Zarr array or group*

Description

Read the .zattrs file associated with a Zarr array or group

Usage

```
write_zarr_attributes(
  zarr_path,
  new.zattrs = list(),
  overwrite = TRUE,
  zarr_version = if (has_metadata_v2) 2L else 3L
)
```

Arguments

zarr_path	A character vector of length 1. This provides the path to a Zarr array or group.
new.zattrs	a list inserted to .zattrs at the path.
overwrite	if TRUE (the default), existing .zattrs elements will be overwritten by new.zattrs.
zarr_version	The version of the Zarr specification to use. If a metadata file already exists, the version will be inferred from the file. Otherwise, the default is 3.

Value

Invisibly, the updated attributes as a named list. This is equivalent to (but faster than) using `read_zarr_attributes()` after writing. If no attributes were present before, this is identical to `new.zattrs`.

Examples

```
z1 <- withr::local_tempdir(fileext = ".zarr")
write_zarr_attributes(z1, list(date = "2025-01-01", author = "Jane Doe"))
```

zarr_overview	<i>Print a summary of a Zarr array</i>
---------------	--

Description

When reading a Zarr array using `read_zarr_array()` it is necessary to know its shape and size. `zarr_overview()` can be used to get a quick overview of the array shape and contents, based on the `.zarray` (Zarr v2) or `zarr.json` (Zarr v3) metadata file each array contains.

Usage

```
zarr_overview(zarr_array_path, s3_client, as_data_frame = FALSE)
```

Arguments

zarr_array_path	A character vector of length 1. This provides the path to a Zarr array or group of arrays. This can either be on a local file system or on S3 storage.
s3_client	A list representing an S3 client. This should be produced by <code>paws.storage::s3()</code> .
as_data_frame	Logical determining whether the Zarr array details should be printed to screen (FALSE) or returned as a <code>data.frame</code> (TRUE) so they can be used computationally.

Details

The function currently prints the following information to the R console:

- array path
- array shape and size
- chunk and size
- the number of chunks
- the datatype of the array
- codec used for data compression (if any)

If given the path to a group of arrays the function will attempt to print the details of all sub-arrays in the group.

Value

If `as_data_frame = FALSE` the function invisible returns `TRUE` if successful. However it is primarily called for the side effect of printing details of the Zarr array(s) to the screen. If `as_data_frame = TRUE` then a `data.frame` containing details of the array is returned.

Examples

```
## Using a local file provided with the package
z1 <- system.file("extdata", "zarr_examples", "row-first",
  "int32.zarr",
  package = "Rarr"
)

## read the entire array
zarr_overview(zarr_array_path = z1)

## using a file on S3 storage

z2 <- "https://noaa-nwm-retro-v2-zarr-pds.s3.amazonaws.com/feature_id/"
zarr_overview(z2)
```

ZarrArray

Deprecated DelayedArray backend functions

Description

[Superseded]

The DelayedArray backend has moved to a dedicated package: the ZarrArray package (<https://github.com/Bioconductor/ZarrArray>).

Usage

```
ZarrArray(...)
```

```
writeZarrArray(...)
```

Arguments

... Passed to the new function in the ZarrArray package.

Examples

```
zarr_path <- system.file(  
  "extdata",  
  "zarr_examples",  
  "column-first",  
  "int32.zarr",  
  package = "Rarr"  
)  
ZarrArray(  
  zarr_path  
)
```

Index

compressors, [2](#), [4](#), [8](#)
create_empty_zarr_array, [3](#)

paws.storage::s3(), [5](#), [6](#), [10](#)

read_zarr_array, [5](#)
read_zarr_array(), [10](#)
read_zarr_attributes, [6](#)

update_zarr_array, [7](#)
update_zarr_array(), [5](#)
use_blosc (compressors), [2](#)
use_bz2 (compressors), [2](#)
use_gzip (compressors), [2](#)
use_lz4 (compressors), [2](#)
use_lzma (compressors), [2](#)
use_zlib (compressors), [2](#)
use_zstd (compressors), [2](#)

write_zarr_array, [8](#)
write_zarr_array(), [5](#)
write_zarr_attributes, [9](#)
writeZarrArray (ZarrArray), [11](#)

zarr_overview, [10](#)
ZarrArray, [11](#)