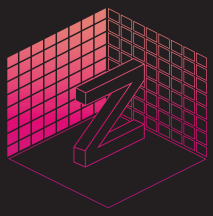

Zarr

Zarr is a container for dense ND-dimensional array data, motivated by the need for a simple, transparent, open, and community-driven format that supports high-throughput distributed I/O on different storage systems.

Zarr data can be stored in any storage system that can be represented as a key-value store, including most commonly POSIX file systems and cloud object storage but also zip files as well as relational and document databases.

Zarr is a Sponsored Project of NumFOCUS, a US 501(c)(3) public charity.

NumFOCUS Sponsored Projects rely on the generous support of corporate sponsors, institutional partners, and individual donors.

For more information on Zarr, including our governance structure and project roadmap, please visit:

https://zarr.dev

APPLICATIONS

Simple and fast serialization of NumPy-like arrays, accessible from languages including Python, C, C++, Rust, Javascript and Java

Multi-scale n-dimensional image storage, e.g. in light and electron microscopy

Geospatial rasters, e.g. following the NetCDF / CF metadata conventions

CURRENT FEATURES

+ Chunk multi-dimensional arrays along any dimension.
+ Store arrays in memory, on disk, inside a Zip file, on S3, etc.
+ Read and write arrays concurrently from multiple threads or processes.
+ Organize arrays into hierarchies via annotatable groups.

PLANNED FEATURES

+ Enable the collection of many small chunk files into superchunks, or “shard” files
+ Improve interoperability with Python data libraries like NumPy, Dask, and Xarray
+ Support new languages like R & MATLAB by wrapping a native Zarr library
+ Simplify the development of extensions by third party developers

NUMFOCUS
OPEN CODE = BETTER SCIENCE

For more information:
info@numfocus.org | +1 (512) 831-2870