1st Conference on Research Data Infrastructure

Poster presentations II

https://doi.org/10.52825/CoRDI.v1i.285

Published: 07 Sept. 2023

Zarr

A Cloud-Optimized Storage for Interactive Access of Large Arrays

Josh Moore[1][https://orcid.org/0000-0003-4028-811X], Susanne Kunis[2][https://orcid.org/0000-0001-6523-7496]

[1] German BioImaging – Gesellschaft für Mikroskopie und Bildanalyse e.V.

[2] Department of Biology/Chemistry and Center for Cellular Nanoanalytics, University Osnabrück, Germany

Abstract. For decades, the sharing of large N-dimensional datasets has posed issues across multiple domains. Interactively accessing terabyte-scale data has previously required significant server resources to properly prepare cropped or down-sampled representations on the fly. Now, a cloud-native chunked format easing this burden has been adopted in the bioimaging domain for standardization. The format — Zarr — is potentially of interest for other consortia and sections of NFDI.

Keywords: FAIR, Community, Bioimaging, Data, Cloud, Format

In an ideal FAIR [1] bioimaging world, the seamless sharing of large image data — dense, often terabyte-scale, N-dimensional arrays — from microscope through to publication and even re-analysis would be possible (Figure 1). The reality, unfortunately, is much less clear due to the lack of a common format for exchange.

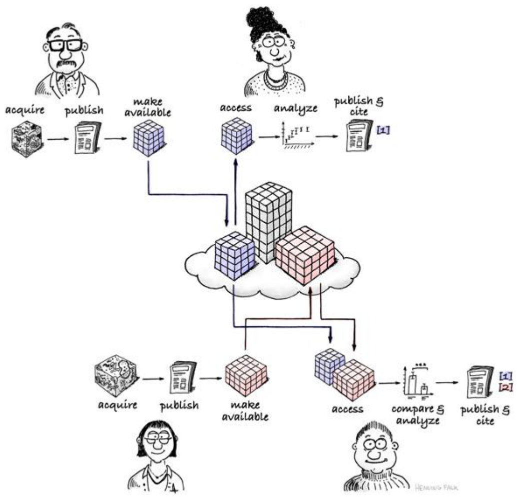

Figure 1. FAIR sharing of data is beneficial for both data producers and consumers. Consumers gain access to interesting datasets that would otherwise be out of reach. Producers get citations to their work when consumers publish their derivative work. OME-Zarr is the technology basis for enabling effective FAIR sharing of large image datasets. "FAIR re-use" by Henning Falk, ©2022 NumFOCUS, is used under a CC BY 4.0 license. [2]

Over 150 different file formats are produced by acquisition systems. Often the data must be duplicated into one or more additional formats for specific applications, increasing storage requirements. Libraries like Bio-Formats [3] can be used to extract the pixel data as well as critical metadata, making it possible to develop server systems like OMERO [4] and the Image Data Resource (IDR) [5]. However, users who would like to download the data are left with the often-complicated translation burden.

To provide a common container while enabling cloud-optimized sharing, the Open Microscopy Environment (OME) began the development of a next-generation file format (NGFF) in 2018 which was subsequently published in 2021 [6]. The basis of this work is Zarr (https://zarr.dev), a format specification for the storage of large N-dimensional typed arrays which avoids certain scalability issues by chunking data into atomic units which can be written and read independently.

Based closely on the HDF5 model [7], Zarr stores hierarchical groups of datasets with arbitrary metadata attached at each level of the structure. Critically, Zarr differs from HDF5 in that rather than using a complex internal binary data structure, a Zarr dataset comprises many individual files so that each chunk or metadata file (written in JSON) can be referenced via predefined, externally stable paths which can be listed by standard file and web browsers (Figure 2). For many of the storage backends, this enables the parallel writing of large image datasets, essential for cluster and cloud-based processing, as well as the viewing of terabytes of data from a static webpage. Implementations exist in several programming languages including C++, Java, JavaScript, Julia, Python, R, and Rust.

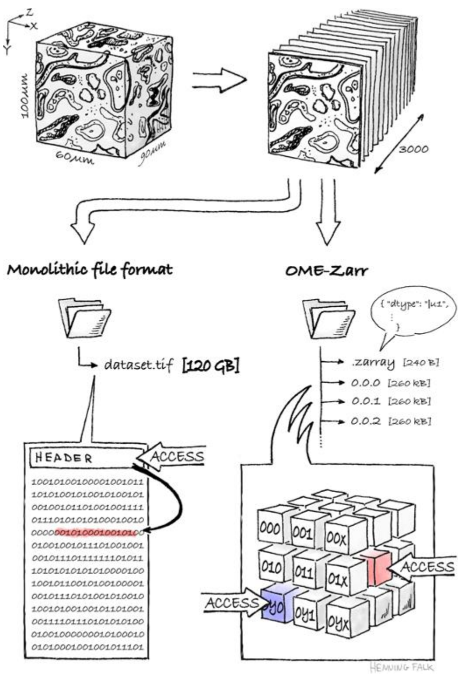

Figure 2. Technical differences between monolithic and chunked file formats. "Monolithic vs. chunked" by Henning Falk, ©2022 NumFOCUS, is used under a CC BY 4.0 license. [2]

By adding OME metadata [8] to the Zarr files, a standard imaging format has been created, supported by multiple visualization tools available with more in development (Figure 3). The flexible metadata structure permits the addition of further schemas like “Recommended Metadata for Biological Images” (REMBI) [9], “Quality Assessment and Reproducibility for Instruments & Images in Light Microscopy” (QUAREP-LiMi) [10], or “Minimum information guidelines for highly multiplexed tissue images” (MITI) [11] in the future. Repositories like the BioImage Archive [12] with the stated goal of accepting all published image data need the scalability afforded by OME-Zarr's serverless capabilities, while the format is also ideal for storing datasets like ZebraHub [13] or the Cell Painting Gallery [14], on institute storage or in Amazon's Open Data program respectively. More information on this growing ecosystem is available in a recent preprint on OME-Zarr [15].

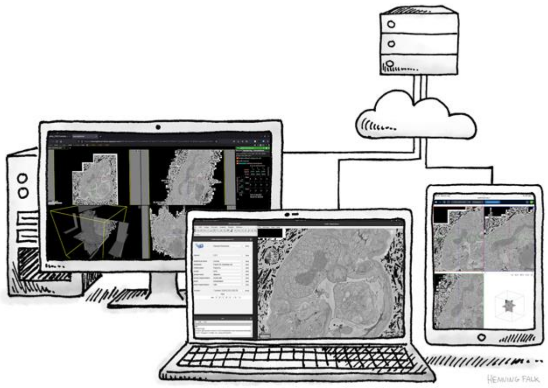

Figure 3. Several visualization tools already support OME-Zarr, including Fiji/Bigdataviewer/MoBIE [16] on the desktop, webKnossos [17] and Neuroglancer [18] on the web. With OME-Zarr, a selection of software tools and devices can access the same datasets from centralized storage. "Multiple clients" by Henning Falk, ©2022 NumFOCUS, is used under a CC BY 4.0 license. [2]

The technical issue which led to the development of Zarr, however, affects other communities. There is significant latency inherent in working with remote data, especially on cloud storage. In the original OME-NGFF paper [6], a benchmark compared the performance of Zarr, TIFF (Tagged Image File Format), and HDF5 files containing the same synthetic data accessed locally, via HTTP, and S3. The benchmark sequentially loaded random, uncompressed chunks to identify the overhead experienced by users visualizing data. It was shown that for monolithic formats like TIFF and HDF5 the overhead of remotely traversing the internal binary structure leads to performance degradation. The pre-computable locations of the Zarr chunks and metadata, in comparison, scale more favorably as latency increases. This is not to say that Zarr is fundamentally preferable to HDF5, but rather that there is a second, cloud-native paradigm with distinct characteristics (Table 1) which need to be considered when sharing large arrays and when developing software infrastructure.

Table 1. Back of the envelop comparison for characteristics of the two primary storage domains – filesystems and object storage, adapted originally from https://www.openio.io/blog/block-file-object-storage-evolution-computer-storage-systems

	Filesystem	**Object storage ("cloud")**
Storage cost	1 EUR/GB	0.01 EUR/GB
Throughput	Gbps	Tbps
Latency	10 µs	1 ms
Modifications	I/O intensive	Immutable

To this end, GUIs and libraries are being updated or created, which multiple institutes are using to migrate their data to OME-Zarr and publish it online. In our recent preprint [15], members of the bioimaging community seek to signal a clear investment in this Zarr-based format, both to industry partners as well as nearby communities. Discussion and the exchange of specifications is already occurring, e.g., with the geospatial community. The Open Geospatial Consortium (OGC) has adopted Zarr as a community standard [19] while NASA divisions like POWER (https://power.larc.nasa.gov/) have migrated their currently NetCDF-based data to Zarr and other divisions are in the process of evaluating the solution. We expect this process to continue as this becomes an agreed upon mechanism for sharing data moving forward.

Task Area 1, “Image (meta)data formats & standardization”, of the NFDI4BIOIMAGE consortium is committed to delivering a Zarr-based format for bioimaging, but as a general purpose, multi-language, extensible yet approachable specification, Zarr could be of interest to other NFDI consortia and the sections.

Competing interests

J.M. is a member of the steering councils of both the OME and Zarr software projects and holds equity in Glencoe Software, a commercial company that builds, delivers, supports, and integrates image data management systems across academic, biotech, and pharmaceutical industries.

Funding

J.M. was supported for work on OME-NGFF by Deutsche Forschungsgemeinschaft (DFG, German Research Foundation) NFDI4BIOIMAGE project 501864659 and by Chan Zuckerberg Initiative DAF grant numbers 2019-207272 and 2022-310144 as well as for work on Zarr by Chan Zuckerberg Initiative DAF grant numbers 2019-207338 and 2021-237467.

Acknowledgement

J.M. would like to thank the international NGFF community (https://ngff.openmicroscopy.org) for the specifications, tools, data resources, and ever lively discussions.

References

1. M. D. Wilkinson et al., “The FAIR Guiding Principles for scientific data management and stewardship,” Sci Data, vol. 3, p. 160018, Mar. 2016, doi: 10.1038/sdata.2016.18.
2. J. Moore, zarr-developers/zarr-illustrations-falk-2022: Zarr illustrations by Henning Falk (August 2022). 2022. doi: 10.5281/zenodo.7037679.
3. M. Linkert et al., “Metadata matters: access to image data in the real world,” J. Cell Biol., vol. 189, no. 5, pp. 777–782, May 2010, doi: 10.1083/jcb.201004104.
4. J.-M. Burel et al., “Publishing and sharing multi-dimensional image data with OMERO,” Mamm. Genome, vol. 26, no. 9–10, pp. 441–447, Oct. 2015, doi: 10.1007/s00335-015-9587-6.
5. E. Williams et al., “The Image Data Resource: A Bioimage Data Integration and Publication Platform,” Nat. Methods, vol. 14, no. 8, pp. 775–781, Aug. 2017, doi: 10.1038/nmeth.4326.
6. J. Moore et al., “OME-NGFF: a next-generation file format for expanding bioimaging data-access strategies,” Nat. Methods, vol. 18, no. 12, pp. 1496–1498, Dec. 2021, doi: 10.1038/s41592-021-01326-w.
7. The HDF Group, Hierarchical Data Format version 5. 1997-2021. [Online]. Available: http://www.hdfgroup.org/HDF5
8. I. G. Goldberg et al., “The Open Microscopy Environment (OME) Data Model and XML file: open tools for informatics and quantitative analysis in biological imaging,” Genome Biol., vol. 6, no. 5, p. R47, May 2005, doi: 10.1186/gb-2005-6-5-r47.

9. U. Sarkans et al., “REMBI: Recommended Metadata for Biological Images-enabling reuse of microscopy data in biology,” Nat. Methods, vol. 18, no. 12, pp. 1418–1422, Dec. 2021, doi: 10.1038/s41592-021-01166-8.
10. G. Nelson et al., “QUAREP-LiMi: A community-driven initiative to establish guidelines for quality assessment and reproducibility for instruments and images in light microscopy,” arXiv [q-bio.OT], Jan. 21, 2021. [Online]. Available: http://arxiv.org/abs/2101.09153
11. D. Schapiro et al., “MITI minimum information guidelines for highly multiplexed tissue images,” Nat. Methods, vol. 19, no. 3, pp. 262–267, Mar. 2022, doi: 10.1038/s41592-022-01415-4.
12. M. Hartley, G. Kleywegt, A. Patwardhan, U. Sarkans, J. R. Swedlow, and A. Brazma, “The BioImage Archive - home of life-sciences microscopy data,” bioRxiv, p. 2021.12.17.473169, Dec. 21, 2021. doi: 10.1101/2021.12.17.473169.
13. M. Lange et al., “Zebrahub – Multimodal Zebrafish Developmental Atlas Reveals the State Transition Dynamics of Late Vertebrate Pluripotent Axial Progenitors,” bioRxiv, p. 2023.03.06.531398, Mar. 07, 2023. doi: 10.1101/2023.03.06.531398.
14. S. N. Chandrasekaran et al., “Three million images and morphological profiles of cells treated with matched chemical and genetic perturbations,” bioRxiv, p. 2022.01.05.475090, Jan. 05, 2022. doi: 10.1101/2022.01.05.475090.
15. J. Moore et al., “OME-Zarr: a cloud-optimized bioimaging file format with international community support,” bioRxiv, Feb. 2023, doi: 10.1101/2023.02.17.528834.
16. C. Pape et al., “MoBIE: a Fiji plugin for sharing and exploration of multi-modal cloud-hosted big image data,” Nat. Methods, vol. 20, no. 4, pp. 475–476, Apr. 2023, doi: 10.1038/s41592-023-01776-4.
17. K. M. Boergens et al., “webKnossos: efficient online 3D data annotation for connectomics,” Nat. Methods, vol. 14, no. 7, pp. 691–694, Jul. 2017, doi: 10.1038/nmeth.4331.
18. J. Maitin-Shepard et al., google/neuroglancer: Zenodo, 2021. doi: 10.5281/ZENODO.5573293.
19. Open Geospatial Consortium. “Zarr Storage Specification 2.0 Community Standard”. https://portal.ogc.org/files/100727 (April 22, 2023)