

Quinteros, J., Carter, J. A., Schaeffer, J., Trabant, C., Pedersen, H. A. (2021): Exploring Approaches for Large Data in Seismology: User and Data Repository Perspectives. - Seismological Research Letters, 92, 3, 1531-1540.

<https://doi.org/10.1785/0220200390>

# 1 Exploring Approaches for Large Data in

## 2 Seismology: User and Data Repository

### 3 Perspectives

4 Authors: Javier Quinteros<sup>1</sup>, Jerry A. Carter<sup>2</sup>, Jonathan Schaeffer<sup>3</sup>, Chad Trabant<sup>2</sup>, Helle A.

5 Pedersen<sup>4</sup>

#### 6 Abstract

7 New data acquisition techniques are generating data at much finer temporal and spatial  
8 resolution compared to traditional seismic experiments. This is a challenge for data centers and  
9 users. As the amount of data potentially flowing into data centers increases by one or two  
10 orders of magnitude, data management challenges are found throughout all stages of the data  
11 flow.

12 The IRIS, RESIF and GEOFON data centers carried out a survey and conducted interviews of  
13 users working with very large datasets to understand their needs and expectations. One of the

---

<sup>1</sup> GFZ German Research Centre for Geosciences, Telegrafenberg, 14473 Potsdam, Germany  
[javier@gfz-potsdam.de](mailto:javier@gfz-potsdam.de)

<sup>2</sup> Incorporated Research institutions for Seismology (IRIS), 9706 4th Ave NE Suite 303, Seattle  
WA 98115, USA

<sup>3</sup> Univ. Grenoble Alpes, Irstea, CNRS, IRD, Météo France, OSUG, 38000 Grenoble, France

<sup>4</sup> Univ. Grenoble Alpes, Univ. Savoie Mont Blanc, CNRS, IRD, Univ. Gustave Eiffel, ISTerre, 38000  
Grenoble, France

14 conclusions is that existing data formats and services are not well suited for users of large  
15 datasets. Data centers are exploring storage solutions, data formats, and data delivery options  
16 to meet large dataset user needs. New approaches will need to be discussed within the  
17 community to establish large dataset standards and best practices, perhaps through  
18 participation of stakeholders and users in discussion groups and forums.

## 19 Introduction

20 New methods of measuring ground motion are significantly reducing the cost of data collection.  
21 Two new types of equipment strongly contribute to this cost reduction. The first is nodal data,  
22 i.e. data from experiments with a very high number of observation points using low-cost  
23 sensors, are now becoming common. As an example, more than 5200 high frequency sensors  
24 were deployed over a period of six months in and around Long Beach, California, USA (e.g. Lin et  
25 al., 2013). The second is distributed acoustic sensing (DAS) technology, using fiber-optic cables,  
26 which is currently being tested and deployed in many locations. As an example, Jousset et al.  
27 (2018) deployed a 15-km-long fiber-optic cable layout on Reykjanes Peninsula, SW-Iceland, with  
28 a distance of 4 meters between sampling points, to study structural features in the Reykjanes  
29 Oblique Rift.

30 Both types of equipment can generate tens to hundreds of terabytes of data at a small fraction  
31 of the cost of using traditional seismometers and geophones for equivalent data volumes. As  
32 the price of collecting data becomes drastically cheaper, a corresponding and dramatic increase  
33 in the volume of data being collected provides a potential scientific bonanza.

34 As dataset sizes increase into the range of tens to hundreds of terabytes however, barriers to  
35 storing, transporting, and processing the data begin to appear. The largest collections of openly

36 available seismological data (e.g., as managed by IRIS, RESIF, and GEOFON) measure the volume  
37 of their decades-spanning repositories in hundreds of terabytes, yet some new (e.g. DAS)  
38 experiments are gathering data volumes well in excess of a hundred terabytes - a significant  
39 fraction of the total data volumes presently stored at the seismological data centers. One might  
40 suggest that just a small fraction of the data that are collected should be preserved as was  
41 suggested decades ago when broadband digital data was first introduced, but tremendous  
42 scientific value has been found in the continuous data that was recorded. Storing these data in  
43 perpetuity is a significant and daunting challenge, as is delivering datasets exceeding a few tens  
44 of terabytes. High-performance/ high-transaction computational resources (HPC/HTC) are  
45 increasingly necessary to process these data and these resources are not typically provided by  
46 the repositories, nor are they co-located with the repositories.

47 This explosion in data volumes is just beginning and the data centers that have traditionally  
48 been the repositories of data for the entire community are not only being asked to host these  
49 data for the wider research community but also to accommodate access to the computational  
50 resources that are needed to process these data sets. Appropriate data management practices  
51 by the data centers must address large data transport, reduced-volume derivative data  
52 products, increased access to HPC/HTC, and data formats that are HPC/HTC-friendly.

53 IRIS, RESIF, and GEOFON, all of which are dedicated to providing free and unrestricted access to  
54 their data holdings, have attempted to identify the needs of the community and to look for  
55 common solutions and best practices for managing very large data sets while maintaining their  
56 traditional data services. We begin by presenting the results of soliciting user expectations for  
57 large data services through a survey of large data providers and users. Next, we describe the  
58 challenges posed by large data from its submission to a data center, to archiving, format

59 considerations, data distribution issues, and finally to the diversity of use cases that data centers  
60 provide and will need to provide to serve the research community. The conclusions of this paper  
61 provide some possible solutions and strategies for dealing with the challenges of  
62 accommodating large data sets. As a community, we must also recognize the environmental  
63 impact of storing, processing and transporting large datasets (e.g. carbon footprint due to  
64 energy consumption, water usage, pollution from backup generators). The strategy that is  
65 eventually implemented by the seismological data centers must take this impact into account  
66 and minimize it when at all possible.

## 67 User Expectations

68 To better serve their users, data centers evaluate the data requests that they receive to improve  
69 their existing services and interact with the scientific research and education communities to  
70 discover future trends. This provides insight into popular data selection parameters, the data  
71 formats requested, and the services that the community uses.

## 72 User Survey

73 A survey was conducted of selected and self-identified large data users; 37 responses were  
74 received. The survey responses indicate that researchers anticipate the following:

- 75 • Large volume use, from 1 to 300+ terabytes, of traditional data (e.g. broadband seismic)  
76 and newer data types such as nodal deployments or DAS.
- 77 • Use of both existing data access tools and mechanisms (web services, Python-based  
78 clients, miniSEED) and newer data access and processing such as HDF5, Zarr, xarray,  
79 especially for larger data sets.

80 The respondents reported the maximum size of the datasets they expect to be working with in  
81 the next 3-5 years. We classified them in three different categories of data volume: 21 small (1-  
82 9 TB), 11 medium (10-50 TB), and 5 large (50+ TB). We consider the small volume data users  
83 relatively well served by current data center capacities, especially as they anticipated using well-  
84 established data formats, processing tools and access mechanisms. Medium volume data users,  
85 however, are observed to split their data requests into many small requests. What we  
86 summarize below are the survey results primarily for the medium and large data users. The  
87 results for four important variables are shown in Figure 1.

88 What data type are the raw data (e.g. broadband, nodal, DAS)?

89 The medium data users primarily identified broadband seismic data, nodal data and a bit of  
90 DAS, while large data users indicated mostly DAS, with some nodal and broadband seismic.

91 From which data centers do you expect to request large data volumes?

92 Users from all categories indicated using a wide variety of data centers, which we interpret to  
93 indicate they will access data from wherever they can get it, with no particular preference.

94 In which data format(s) would you prefer to work with large data (miniSEED, SAC, PH5, ASDF,  
95 HDF5, etc.)?

96 Both medium and large data users anticipated a use of a variety of data formats, but primarily  
97 miniSEED, HDF5 (PH5 and ASDF), with a few cases of Zarr (see Data and Resources). However,  
98 for users planning to work with 20+ TB, the option of miniSEED reduces considerably in  
99 comparison to HDF5-based formats or other less standard formats.

100 Would you process the data with standard codes? With your own code? Using third party  
101 frameworks and libraries (e.g. ObsPy)?

102 Users from all categories indicated the most common codes they expect to use would mostly be  
103 their own code, supported by ObsPy (Beyreuther et al., 2010) or MATLAB.

104 Which programming language(s) do you, or would you expect to, use for data processing and  
105 analysis?

106 Python is the most indicated programming language of choice, followed by MATLAB, C/C++,  
107 Fortran and Julia.

108 Where would you prefer the data to be delivered: your own compute infrastructure, a cloud  
109 system, an HPC center?

110 Many medium and large users indicated that they would like to transfer data to their own  
111 computer infrastructure or their institutional infrastructure. Many also indicated a preference to  
112 deliver data to an HPC center one of whom anticipated relying on commercial cloud providers.

113 Would you be willing to pay for the data processing resources (e.g. compute) if the data are  
114 available within a commercial cloud or other environment that is not free to use?

115 Many users are willing to pay (some dependent on grant allowance) for processing. However, as  
116 in previous questions, among users working with 20+ TB there seems to be a clear trend in favor  
117 of paying. Some pointed out that it should be an option to transfer the data to a location where  
118 they would not be required to pay for processing.

119 Please briefly summarize your current large data processing workflow, trouble points  
120 identified/foreseen, and how it can be improved.

121 Most users indicated that they use largely ad hoc data processing workflows, with a few  
122 identifying the need for parallel processing and processing-ready data access such as from HDF5  
123 containers. The common challenge for many is efficient access to large volume data, either due  
124 to access interface issues or limited transfer capability.

125 Which software would you use to access/download large data volumes (libraries, applications,  
126 etc.)?

127 Users of all scales report they expect to use the existing tools and mechanisms of web services,  
128 ObsPy, or directly in their own codes (in Python or Julia). A few report their anticipated use of  
129 larger platforms such as Pangeo, direct access to object storage (e.g. S3), specialised libraries  
130 (e.g. xarray), and advanced data containers such as Zarr.

131 Other comments

132 When asked for other comments or their vision for large seismic data processing in the future,  
133 users identified the need to develop access and transfer mechanisms for large volumes of data.  
134 Additionally, they noted the need for advanced data formats appropriate for large data, the  
135 desire for derived (reduced volume) data products for easier use, their preference for compute  
136 resources in the same system as the data (to avoid transfer), and cloud ready seismic processing  
137 software. Finally, the users expressed their desire for continued collaboration between  
138 international seismological data centers so that researchers can discover and access large  
139 volumes of data wherever it is available.

140 [Diversity of Data Requests](#)

141 From the survey and from the experience of the IRIS, GEOFON, and RESIF data centers, requests  
142 for medium and large data sets span a range of time and spatial scales. On one end of the  
143 spectrum are studies that require hours/days of high sample-rate data collected from densely  
144 spaced sensors, and at the other end of the spectrum are studies that require decades of broad-  
145 band data from stations distributed around the globe.

146 The survey clearly demonstrates that the DAS experiments represent the largest volume  
147 datasets; however, both the survey and data center experience show that Medium and Large



148 data volumes from traditional seismic stations are increasingly in demand. Two typical uses of  
149 such Medium and Large datasets include studies based on cross correlation techniques and  
150 studies using machine learning methods (ML). In such cases, large datasets are created from  
151 years-long time series from a large number of individual seismic stations either at global,  
152 regional or local scales and the data needed for a comprehensive study may even be distributed  
153 across several data centers on different continents. In contrast, requests for node data or other  
154 experimental data with high frequency sampling are often based on data from a single data  
155 center. Users may want access to the entire dataset or to a specific time or spatial slice. DAS  
156 datasets can be much larger than anything that traditional data centers supporting the scientific  
157 research and education communities have experience in managing. Like nodal data, researchers  
158 may want access to the entire dataset or to subsets of the full dataset; they may also want to  
159 process those data without transporting them or need to have them delivered to an HPC center  
160 rather than to their home institutions. Because of these processing needs, data centers should  
161 be aware of the data formats that are efficient for high-performance computation and be able  
162 to deliver data in some of those formats.

163 The management and organization of data and user services need to be flexible enough to  
164 handle not only the requests for medium and large datasets, but also for the diversity of  
165 standard data requests for a small number of seismic stations and short time periods (small  
166 datasets). Data centers need to encourage and enhance services associated with data requests  
167 using criteria such as station location, instrument type, sampling rate, data quality, data  
168 repository location, or more advanced parameters associated with earthquake parameters,  
169 propagation path, local geology at the seismic station, etc. MUSTANG (Casey et al., 2018) and

170 WFCatalog (Trani et al, 2017) are some examples of services that allow users to pre-select data  
171 based on quality metrics.

## 172 The Challenges of Large Data

173 The seismological community has a long tradition of collaborative work between data centers  
174 and researchers. In Europe data centers collaborate through Observatories & Research Facilities  
175 for European Seismology (ORFEUS), while global collaboration in the seismological community is  
176 through the International Federation of Digital Seismograph Networks (FDSN). One of the most  
177 important achievements for the discipline is the establishment of well-defined standards under  
178 the coordination of the FDSN. These include not only the data formats, but also detailed  
179 specifications for providing data/metadata services and how Digital Object Identifiers (DOIs)  
180 should be used to identify seismic networks. Standards approved by the global community  
181 within the FDSN allow seismological data centers to support the Findable, Accessible,  
182 Interoperable, and Repeatable (FAIR) data principles (Wilkinson et al., 2016). This global  
183 collaboration has contributed to a greater sharing of data among all seismologists and is a  
184 model for other disciplines.

185 Specifications regarding formats and services were designed in a totally different landscape of  
186 operational limitations and data usage compared to what the community will face from now on.  
187 As soon as the amount of data increases one or two orders of magnitude, technical problems  
188 are expected. From the data center perspective, large data management challenges are found in  
189 the following stages of the data flow: (a) data submission, (b) data archival, (c) creation of  
190 proper metadata, (d) data distribution, and (e) data usage (Figure 2).

## 191 [Data Submission](#)

192 In the cases of Large-N or DAS experiments the volume of data submitted to a data center could  
193 be hundreds of terabytes. Here, it is important to make a distinction between temporary and  
194 permanent experiments. For temporary experiments, data are usually acquired in the field and  
195 either transmitted or physically transported to the data center. Field data storage devices can  
196 be mounted on the data center's network and the data ingested into the system. In the case of  
197 permanent deployments, data are usually transmitted to the data center in real-time by means  
198 of protocols like SeedLink. However, for high volume data such as from DAS experiments, real-  
199 time transmission of the full-resolution data may not be practical. One option for DAS  
200 installations is to pre-process the data at the station to reduce the data volume before  
201 transmitting them to a data center. Raw data would still be available for later delivery to a data  
202 center or, should no data center be capable of handling the volume, a suitable archive found by  
203 the operator of the DAS installation.

## 204 [Data Archive Planning](#)

205 Designing an efficient data management system for seismological data requires consideration of  
206 several aspects: the storage hardware, the data format(s), and the way data will be accessed.  
207 These three factors are connected; storage system performance varies depending on user  
208 access requirements, which are supported by the format of the data. Requests that require  
209 gathering small snippets of data from many files require higher performance storage systems  
210 than requests for data in a single file.  
211 The internal data management of seismological data centers is not specified by FDSN; data  
212 centers are free to manage their data as they see fit and store the data they have in their

213 repositories and archives in whatever format suits their use cases. The data provisioning system,  
214 fdsnws-dataselect (see Data and Resources), must deliver data in the miniSEED format, but the  
215 data centers are free to use different data formats to archive data and perform the conversion  
216 on-the-fly before sending them to users.

217 To date, the miniSEED format has been very popular as a storage format among data centers  
218 because it is highly compatible with the software ecosystem in the community. The main  
219 advantages are that there is no need to convert on-the-fly when sending most data to users and  
220 that datasets can be formed by a concatenation of miniSEED records that can be streamed as  
221 they are retrieved from the storage media. The ability to stream data directly avoids the need to  
222 stage the dataset on the server side before sending it. The disadvantage is that miniSEED has no  
223 inherent indexing capability; users typically use the file system as a means to organize and  
224 locate data, which may require a huge number of files for medium and large datasets.

225 For medium and large datasets, the HDF5 format is popular among users of HPC. Two popular  
226 HDF5-based formats in our community are PH5 (Hess et al., 2018) and ASDF (Krischer et al.,  
227 2016). Initial experimentation has indicated that a reduction in storage space is possible using  
228 PH5, likely due to the longer data segments being compressed. A redesign of the PH5 format to  
229 make it more versatile is expected to be completed in 2021. The new format is expected to be  
230 usable not only for nodal data, but also for DAS, Magnetotelluric, geodetic, and other data  
231 types. A generic HDF5-based solution could have some advantages not only for data centers but  
232 also for users needing to process their large datasets in an HPC-friendly format. The main  
233 disadvantage for data centers is that HDF5 formatted data cannot currently be streamed on the  
234 fly as they are read by the data provisioning systems (see Data Distribution). However, there is a

235 rich ecosystem of interesting open source tools supporting HDF5 such as the THREDDS Data  
236 Server (Unidata, 2020) or HSDS distribution services (see Data and Resources).

237 Other file formats were mentioned in the user’s survey, but data center experience with these  
238 formats is limited. The desirable features of formats like Zarr, which allow dynamic data  
239 chunking, haven’t been thoroughly tested in the community. In Table 1 we summarize some of  
240 the file formats and their characteristics. The different data format features must be explored to  
241 optimize the trade-offs between efficient data compression, fast data extraction, and streaming.  
242 Increases of volume also imply new strategies for choice of physical storage. If data centers are  
243 expected to increase their archive sizes by at least one order of magnitude (or even more), it  
244 might not be feasible to keep all of the data online (Figure 3). Offline storage should probably be  
245 considered inside the data center or provided by external services. This will raise the storage  
246 capacity but data would have to be staged on demand, and with limited size and retention time.  
247 Considering the costs, one should consider that the long-term archival of data implies not only  
248 the cost of the storage units, but also the backup strategy, the renewal of the hardware (e.g.  
249 hard drives) every 5-7 years, and the regular procedures to ensure that data are still readable.

250 One way to optimize data storage is to utilize tiered storage strategies; keep data that are rarely  
251 accessed in slower, cheaper systems and data that are requested frequently on systems that  
252 provide rapid access. Managing how data are tiered could potentially be dynamic, based on past  
253 access patterns. However, these patterns are not easily discerned.

254 Some data centers have already tested commercial cloud services as a storage system to  
255 evaluate the feasibility of migrating their archives to the cloud. Despite the technical advantages  
256 of this approach, there are financial drawbacks that deserve serious consideration, particularly  
257 for large data centers that hold petabytes of data in their repositories, that require substantial

258 compute resources to manage their data, and that deliver petabytes of data each year to their  
259 users.

260 An important consideration regarding choice of storage system is whether the data will be  
261 accessed directly for processing in addition to supporting general extraction systems. Direct  
262 data access adds a new dimension to the usage patterns and presents a challenge for managing  
263 access permissions, request rates, response times, etc. Direct data access for processing should  
264 not be allowed to deteriorate services associated with standard data requests.

## 265 [Data Description](#)

266 There exist comprehensive metadata standards in seismology that are well suited for large  
267 nodal data sets or other classical seismic data sets. However, newer, emerging data types such  
268 as DAS are not as well handled by these standards where the concepts of a network, station,  
269 location and a channel are different and the fundamental instrument responses are new.  
270 Therefore, one of the most difficult tasks for data centers attempting to identify and implement  
271 standard services is to standardize DAS metadata through the definition of new metadata  
272 schemas. StationXML, the international standard for seismic metadata, is not a good fit with the  
273 technical features of DAS interrogators and the generated datasets. There is a need to tackle  
274 this in a smart way in order to make these data available and interoperable with all popular  
275 tools used within the community, which expect typical seismic waveforms with some minimum  
276 set of requirements on their metadata, and the possibility of mixing different data types.  
277 A non-technical (and still pending) issue is how to name DAS data sampling points. If the first  
278 approach to standardize and distribute DAS experiment data is to generate derived products as  
279 suggested by many users, there are some details to be considered. Should DAS experiments be  
280 registered by the FDSN in order to get a formal network code? If this is not the case, how could

281 a user identify or select the required sampling points? To have a network code registered would  
282 allow the dataset to be uniquely distinguished from other experiments. This would also allow  
283 one to formally assign a DOI and the proper citation to the experiment, as agreed and expected  
284 in the community.

285 At levels finer than the network, should each sampling point be considered a station, a location  
286 or a channel? If we consider that each fiber could have thousands of sampling points with a  
287 spatial separation as small as of 2 to 4 meters and a length of tens of kilometers, a station  
288 designation seems to be the best fit. As stations have a unique position, we cannot consider  
289 such a deployment to be a single-station-multiple-channels configuration. We note that the ISC  
290 applies a 1-km station rule for the International Registry of Stations that is generally followed by  
291 the community. Quoting ISC: “Because of the need for accurate station positions for hypocenter  
292 location programs, a new international code is assigned if a station is moved more than one (1)  
293 kilometer from the previous location”. Although having locations codes more than a kilometer  
294 away from the station does not technically violate the rule about stations that move, we believe  
295 that it would be extremely confusing to have channels with a location up to 50 kilometers away  
296 from their associated station. We also note that some large array deployments (like the  
297 Norwegian Seismic Array Network, see Data and Resources) identify each array element with  
298 different station names.

299 A reasonable approach would be to define a station for each sampling point, similar to the  
300 situation in nodal experiments where each node is typically assigned a station code. These  
301 naming conventions will be fundamental for interoperability between data centers, as well as a  
302 homogeneous definition of the subsets independent of who is archiving the data.

303 As described above, simple information such as latitude, longitude and orientation are already a  
304 challenge, whose solution is not trivial and could potentially make a difference in the results if  
305 data should be automatically processed. It seems unavoidable to perform some post-processing  
306 of the data to calculate the exact location of each sampling point. This should be based on the  
307 GPS information from the interrogator and the detection of an active source (e.g. “tap test”)  
308 during the deployment phase. Standard and generic codes, benchmarks and guidelines provided  
309 to/by the community would be desirable, and information about the postprocessing needs to  
310 integrate the metadata.

311 Some of these topics are being discussed within community user groups, like the Distributed  
312 Acoustic Sensing Research Coordination Network (DAS RCN, see Data and Resources) and  
313 hopefully some proposals to standardize the description of DAS experiments will emerge. The  
314 needs and related challenges from the irruption of new technologies, like DAS interrogators,  
315 could be the triggering factor to discuss the long-term evolution of our current standards (e.g.  
316 StationXML) and to potentially explore more generic formats like SensorML, which is widely  
317 used in Geophysics and generic enough to encompass all that StationXML can express.

### 318 [Data Distribution](#)

319 Currently, the FDSN-standard suite of web services (fdsnws-station and fdsnws-dataselect) are  
320 used by many researchers to discover and download data. From April to September 2020, 90%  
321 of the fdsnws-dataselect requests received delivered less than 5 MB of data (Figure 4). The  
322 miniSEED format along with the fdsnws-dataselect webservice are well designed for this  
323 purpose. However, users may make many small requests to optimize their data retrieval success  
324 and assemble the complete dataset that they desire on their computers. The size of the  
325 datasets needed by the community is probably larger than is observed in Figure 4. It is still



326 unclear what the download pattern for Large-N or DAS data would be, but downloads of entire  
327 datasets are expected.

328 The distribution of large data sets (50+ TB) is challenging due to bandwidth constraints. As  
329 mentioned previously, one of the main results of the survey we conducted was that the full  
330 resolution data is not needed for most studies, particularly for DAS. Instead, a set of products  
331 derived from the original data would satisfy most users. The derived data products are often a  
332 significantly reduced volume making them more feasible for archiving and distribution by  
333 existing data centers and systems. While the derivative products are useful in reducing the  
334 burden of transferring bulky data, the need to transfer large volumes is not eliminated for all  
335 users.

336 Another significant challenge is the lack of available services to handle the transmission of large  
337 volumes. The current data transfer standard in global broadband seismology (fdsnws-  
338 dataselect), defines a synchronous web service interface that accepts arbitrary data selections  
339 and is expected to begin returning data with little or no delay. It has become a highly successful  
340 mechanism, but although allowed, requesting large data volumes with this service has the risk  
341 of causing connection timeouts. It is difficult to efficiently resume transfers following broken  
342 connections, which are increasingly likely with very large volumes. Also, maintaining acceptable  
343 performance can require significant software and system engineering by the data centers.

344 One tool, IRIS's ROVER (see Data and Resources), has overcome the issue of dropped  
345 connections by using a smart client to manage the synchronous data transfer in miniSEED, and  
346 constructing the final dataset on the client side. This approach requires a relatively complex  
347 client, but is a robust and efficient method of requesting large datasets.

348 Future distribution capabilities worth exploration are the use of GridFTP (Allcock et al., 2005), or  
349 rsync (see Data and Resources) for efficient bulk file transfer. But these methods can only be  
350 used with existing files and not arbitrary data selections. Another possibility is the extension of  
351 web service interfaces like fdsnws-dataselect to allow asynchronous, or batch, data transfer to a  
352 site (e.g. HPC center) identified by the requestor. In other cases and depending on the archive  
353 format, some middleware solutions such as THREDDS Data Server (Unidata, 2020) or HSDS (see  
354 Data and Resources) could potentially be useful solutions for sub-selecting and transforming  
355 large data volumes.

### 356 [Using the Data](#)

357 Considering the classical approach of data usage by seismologists, some important  
358 requirements have been identified in the previous sections, in particular the need for standard  
359 data and metadata formats, and correct and complete metadata. Presently, only miniSEED is an  
360 FDSN standard within the community, but this format is inefficient and cumbersome for  
361 processing large datasets. Non-FDSN standard formats (e.g. HDF5 based) are presently being  
362 used by researchers working with medium and large datasets (20+TB) and they would like to get  
363 data in this format from the data center provisioning systems. The definition of an additional  
364 FDSN standard aimed at large datasets would spur the development of community processing  
365 tools, and would greatly further the cooperation and sharing of resources between data  
366 centers.

367 The user's need to process data in a remote environment (e.g. cloud computing, containers),  
368 means that other factors should also be considered by data centers. For instance, even if a data  
369 center can provide a container or Jupyter notebook (see Data and Resources) with direct  
370 processing of the data, solutions need to be developed to limit restricted dataset access to

371 authorized users. This is not a minor issue, as most of these large datasets go under an embargo  
372 period of some years before being opened to the public. This means that depending on the  
373 user, only part of the archive should be accessible to any given user from the computing  
374 interfaces. A similar issue is related to the data organization (format and structure). The data  
375 center can expose (or let the user access) what exists in the storage system so the user will in  
376 principle be able to use only the storage format. It will be difficult (and in some cases  
377 impossible) to stage a new dataset in another format to be accessed by the user. Therefore, the  
378 data format and organization needs to be known by the user. As each data center is free to  
379 organize the data in whatever way is optimal to them, data processing software cannot easily be  
380 used across different data center computational facilities.

381 The transparent federation of data is probably the biggest challenge for data centers because it  
382 includes not only all the improvements and developments mentioned in this article, but also the  
383 edge case in which multiple large datasets are requested from different data centers. In such a  
384 case, even running the code on top of the data is not a solution. That would only save the  
385 transfer from that data center, but there is still the problem of requesting all the other data.  
386 Some public initiatives exist in the United States and Europe to propose solutions to the  
387 problem of federated data (e.g. XSEDE/Jetstream and the European Open Science Cloud).  
388 However, there are still no solutions available that cover all the aspects mentioned above.

## 389 Conclusions

390 Despite present initiatives trying to tackle the problem of big seismological datasets and how  
391 they can be archived, delivered, and processed, data centers are facing a situation similar to  
392 that from the mid-80's. Namely, the amount of data generated was much bigger than the data

393 centers could safely and sustainably archive. At that moment, the temporary solution was to  
394 keep only time windows of interest (e.g. waveforms related to an event) and discard the rest of  
395 the data after some time. What has been learned from this experience is that discarding data is  
396 not the best option to consider, as illustrated by present day research in data mining, new types  
397 of seismic signals, and imaging techniques using seismic noise.

398 With the present increase of data volumes come additional problems associated with the  
399 internal organization of data within the data centers and meeting the need for data formats that  
400 are well adapted for storing and slicing big datasets. Along with these problems, there is a need  
401 for adapting present data request tools or developing new ones.

402 At present, the IRIS, GEOFON and RESIF data centers recognize that it is not possible for them to  
403 permanently archive very large datasets such as DAS due to prohibitive costs of long-term  
404 archiving. Safe storage (involving multiple copies) of such large datasets available online is  
405 presently not possible either. Therefore, the seismological data centers should explore a  
406 strategy for delivering datasets through automatic asynchronous services to designated  
407 destinations such as computing facilities where data pre-processing or processing could occur.  
408 For very large datasets, a common strategy is needed for storing only standard data products  
409 such as spatially and/or temporally down-sampled data, and/or small windows of highly  
410 sampled data.

411 Ideally, in the long term, international standards and products should be developed within the  
412 framework of the FDSN. As a preparatory phase, we suggest to undertake the following actions:

- 413 • Encourage a wide and continuous international collaboration on DAS data. This would  
414 include dedicated workshops and special sessions in conferences involving data centers  
415 and scientific users who could define a limited set of standard products for DAS data.

416 One natural forum would be a strong international presence in the Distributed Acoustic  
417 Sensing Research Coordination Network (DAS RCN), which involves both scientific users  
418 and data center operators/managers.

- 419 • Extend the present discussion between our three data centers to other interested data  
420 centers. One option is to set up dedicated technical workshops on very focused subjects,  
421 including solutions to keep the raw data safely for future use. Our community should in  
422 these workshops interact with data centers from other disciplines to understand  
423 similarities and differences of use cases and technical constraints, to benefit from past  
424 experience over a broad range of scientific areas, and to explore possible common  
425 technical solutions.

426 This preparatory work hopefully will lead to new community standards and user services that  
427 are well adapted to large datasets.

## 428 Data and Resources

429 The SEED Manual can be downloaded from [http://www.fdsn.org/pdf/SEEDManual\\_V2.4.pdf](http://www.fdsn.org/pdf/SEEDManual_V2.4.pdf) .

430 Information about HSDS can be read from <https://www.hdfgroup.org/solutions/highly-scalable->

431 [data-service-hsds/](https://www.hdfgroup.org/solutions/highly-scalable-data-service-hsds/) . Documentation of Zarr can be found at <https://zarr.readthedocs.io/> .

432 Information about the DAS Research Coordination Network (DAS RCN) can be found at

433 [https://www.iris.edu/hg/initiatives/das\\_rcn](https://www.iris.edu/hg/initiatives/das_rcn) . Specifications of the FDSN web services can be

434 found at <http://fdsn.org/webservices/FDSN-WS-Specifications-1.2.pdf> . Information on the

435 Norwegian Seismic Array Network (NO) can be found at

436 <http://www.fdsn.org/networks/detail/NO/> . Details of the International Registry of

437 Seismographic Stations can be read from <http://www.isc.ac.uk/registries/registration/> .

438 Documentation of ROVER can be found at <https://iris-edu.github.io/rover/> . Rsync information  
439 available at <https://rsync.samba.org/> . Information about Jupyter Notebooks is available at  
440 <https://jupyter.org/> . Plots were made using Matplotlib (Hunter, 2007).

441 The Portable Array Seismic Studies of the Continental Lithosphere facility (PASSCAL) is a facility  
442 of the IRIS consortium that provides instrumentation for National Science Foundation,  
443 Department of Energy, and otherwise funded seismological experiments around the world.

## 444 Acknowledgments

445 The authors thank the users and data providers for their continuous feedback, and in particular  
446 the users who participated in the survey. Javier Quinteros collaborated in this work as part of  
447 the EOSC-Pillar project, which has received funding from the European Union's Horizon 2020  
448 research and innovation programme under grant agreement No 857650, as well as the RISE  
449 project, also supported by the European Union's Horizon 2020 research and innovation  
450 programme under grant agreement No 821115. GEOFON operates the GFZ seismological data  
451 archive, where own and third party data are curated alongside with data from passive GIPP  
452 experiments. The seismological facility for the advancement of geoscience (SAGE) is operated by  
453 IRIS on behalf of the National Science Foundation (NSF) through Award EAR-1851048. The RESIF  
454 Data center is operated on behalf of the RESIF Consortium, a national Research Infrastructure  
455 coordinated by CNRS and which receives funding, personnel and other from all the Consortium  
456 members. Résif-SI is also supported by the System of Observation and Experimentation for  
457 Research and the Environment (SOERE) and by the Ministry of Ecological Transition. Résif-SI has  
458 benefited from the Résif-CORE project (11-EQPX-0040), funded by the national French funding  
459 program Investissements d'Avenir and managed by the French National Research Agency (ANR).

## 460 References

- 461 Allcock, W., J. Bresnahan, R. Kettimuthu, M. Link (2005). The Globus Striped GridFTP Framework  
462 and Server, SC '05: Proceedings of the 2005 ACM/IEEE Conference on Supercomputing, doi:  
463 10.1109/SC.2005.72.
- 464 Beyreuther, M., R. Barsch, L. Krischer, T. Megies, Y. Behr, J. Wassermann (2010). ObsPy: A  
465 Python Toolbox for Seismology, *Seismol. Res. Lett.*, doi: 10.1785/gssrl.81.3.530.
- 466 Casey, R., M. E. Templeton, G. Sharer, L. Keyson, B. R. Weertman, T. Ahern (2018). Assuring the  
467 Quality of IRIS Data with MUSTANG, *Seismol. Res. Lett.*, doi: 10.1785/0220170191.
- 468 Godoy, W. F., N. Podhorszki, R. Wang, C. Atkins, G. Eisenhauer, J. Gu, P. Davis, J. Choi, K.  
469 Germaschewski, K. Huck, A. Huebl, M. Kim, J. Kress, T. Kurc, Q. Liu, J. Logan, K. Mehta, G.  
470 Ostrouchov, M. Parashar, F. Poeschel, D. Pugmire, E. Suchyta, K. Takahashi, N. Thompson, S.  
471 Tsutsumi, L. Wan, M. Wolf, K. Wu, S. Klasky (2020). ADIOS 2: The Adaptable Input Output  
472 System. A framework for high-performance data management, *SoftwareX*, doi:  
473 10.1016/j.softx.2020.100561.
- 474 Hess, D., N. Falco, rsdeazevedo, damhuonglan, K. Jacobs (2018). PIC-IRIS/PH5: v4.1.2 (Version  
475 v4.1.2). Zenodo. doi: 10.5281/zenodo.1284569.
- 476 Hoyer, S., J. Hamman (2017). xarray: N-D labeled Arrays and Datasets in Python, *J. Open Res.*  
477 *Software*, doi: 10.5334/jors.148.
- 478 Hunter, J. D. (2007). Matplotlib: A 2D graphics environment, *Comput. Sci. Eng.* 9, no. 3, 90-95.
- 479 Jousset, P., T. Reinsch, T. Ryberg, et al. (2018). Dynamic strain determination using fibre-optic  
480 cables allows imaging of seismological and structural features, *Nat Commun*, doi:  
481 10.1038/s41467-018-04860-y

482 Krischer, L., J. Smith, W. Lei, M. Lefebvre, Y. Ruan, E. Sales de Andrade, N. Podhorszki, E. Bozdağ,  
483 J. Tromp (2016). An Adaptable Seismic Data Format, *Geophys. J. Int.*, doi:  
484 10.1093/gji/ggw319.

485 Lin, F.-C., D. Li, R. W. Clayton, D. Hollis (2013). High-resolution 3D shallow crustal structure in  
486 Long Beach, California: Application of ambient noise tomography on a dense seismic array,  
487 *Geophysics*, 78 (4), Q45–Q56, doi: 10.1190/GEO2012-0453.1

488 SEED (2012) SEED Reference Manual. Standard for the exchange of earthquake data, SEED  
489 format version 2.4. International Federation of Digital Seismograph Networks Incorporated  
490 Research Institutions for Seismology (IRIS), USGS.

491 Trani, L., M. Koymans, M. Atkinson, R. Sleeman, R. Filgueira (2017). WFCatalog: A catalogue for  
492 seismological waveform data, *Comput. Geosci.*, doi: 10.1016/j.cageo.2017.06.008.

493 Unidata (2020). THREDDS Data Server (TDS) version 4.6.2 [software]. Boulder, CO:  
494 UCAR/Unidata. doi: 10.5065/D6N014KG.

495 Wilkinson, M., M. Dumontier, I. Aalbersberg, et al. (2016). The FAIR Guiding Principles for  
496 scientific data management and stewardship, *Sci. Data*, doi: 10.1038/sdata.2016.18.  
497



498 Full mailing address for each author

499 Javier Quinteros - GFZ German Research Centre for Geosciences, Telegrafenberg, 14473

500 Potsdam, Germany javier@gfz-potsdam.de

501 Jerry A. Carter - Incorporated Research institutions for Seismology (IRIS), 9706 4th Ave NE Suite

502 303, Seattle WA 98115, USA

503 Jonathan Schaeffer - Univ. Grenoble Alpes, Irstea, CNRS, IRD, Météo France, OSUG, 38000

504 Grenoble, France

505 Chad Trabant - Incorporated Research institutions for Seismology (IRIS), 9706 4th Ave NE Suite

506 303, Seattle WA 98115, USA

507 Helle A. Pedersen - Univ. Grenoble Alpes, Univ. Savoie Mont Blanc, CNRS, IRD, Univ. Gustave

508 Eiffel, ISTerre, 38000 Grenoble, France

509

510

511 Table 1: Summary of data formats mentioned in the survey and some of their relevant  
 512 advantages and disadvantages.

Format	Comments regarding use in large data management
miniSEED SEED (2012)	miniSEED is an international seismological standard for data exchange that is also commonly used for archiving. It was designed for independent, single time series, usually very small chunk sizes, and not for processing.
PH5 (Hess et al., 2018)	PH5 is an HDF5 based format. A rich toolbox has been provided by PASSCAL but data subsetting is not possible without high-level translation and streaming the base format is not supported.
ASDF (Krischer et al., 2018)	ASDF is an HDF5 based format. Data subsetting is not possible without high-level translation and streaming the base format is not supported.
Zarr ZARR (2020)	Zarr is a format and a python library to manage numerical arrays. Data subsetting is allowed and it is a good design for HPC. Zarr has not been used much in the geophysical community and is only for python.
ADIOS2 (Godoy et al., 2020)	ADIOS2 is a framework for data IO management; it is not broadly known or used in seismology.

513

514

515 Figure Captions

516

517 Figure 1: Different aspects of users' data workflows by size of datasets to be requested. Counts  
518 at the base of bars are the number of respondents. Smaller datasets tend to be requested in  
519 miniSEED format, but HDF5-based formats are preferred for larger datasets. Requests to get  
520 nodal experiment data will generate medium size datasets, while users plan to get very large  
521 datasets from DAS experiments. Python is the preferred programming independent of the  
522 dataset size. Small datasets are expected to be staged in the user infrastructure, while larger  
523 datasets in some HPC facility. Only a few users mentioned their will to stage data in a cloud  
524 service.

525

526 Figure 2: Schematic view of the different parts of the data flow in a seismological data center.  
527 The inputs are (a) data and (c) data description (metadata), submitted by the producer. Both  
528 inputs are managed in separate workflows to be hosted in (b) data archive and (c') metadata  
529 database. Data distribution services (d) need access to both repositories in order to handle the  
530 (e) final user's requests. Challenges regarding large data in a, b, c, d and e are described in this  
531 paper.


532

533 Figure 3: Evolution of the data archived yearly at each data center (green, yellow, and blue  
534 curves). The error bars on the right show the expected size range of a single dataset for a Large-  
535 N (blue) and a DAS (red) experiment. It can be seen that a single large dataset could be  
536 equivalent in size to a whole year of the typical datasets currently being archived.

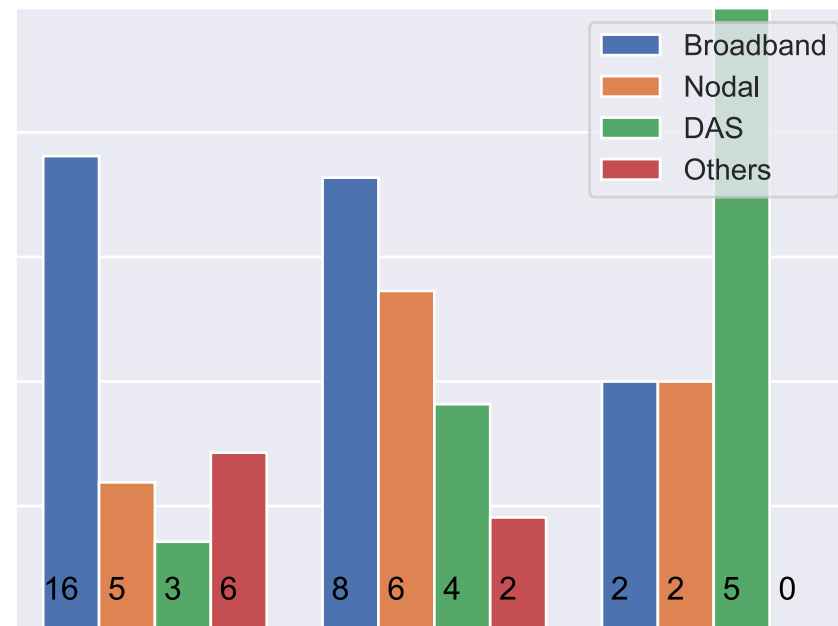
537

538 Figure 4: Percentage of requests per request size and data center. Statistics have been  
539 calculated from April 2020 to September 2020. 90% of the requests are less than 5 MB. Note  
540 however that some users request intermediate size volumes through thousands of small data  
541 requests. The dataset sizes needed by users could be much larger than what is shown here.  
542

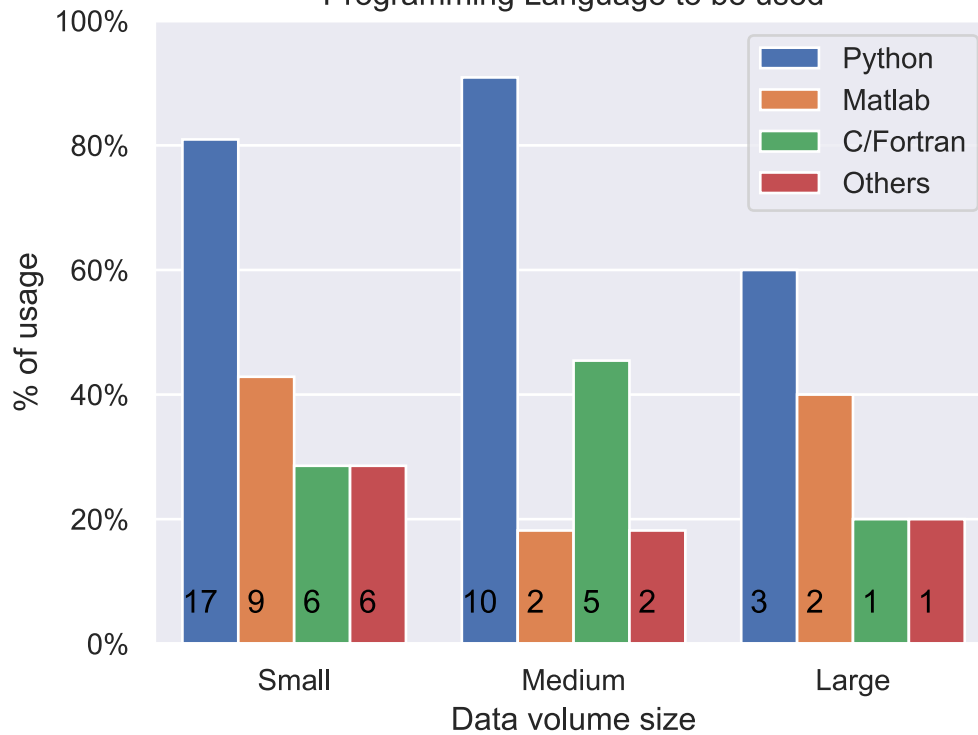
Data format



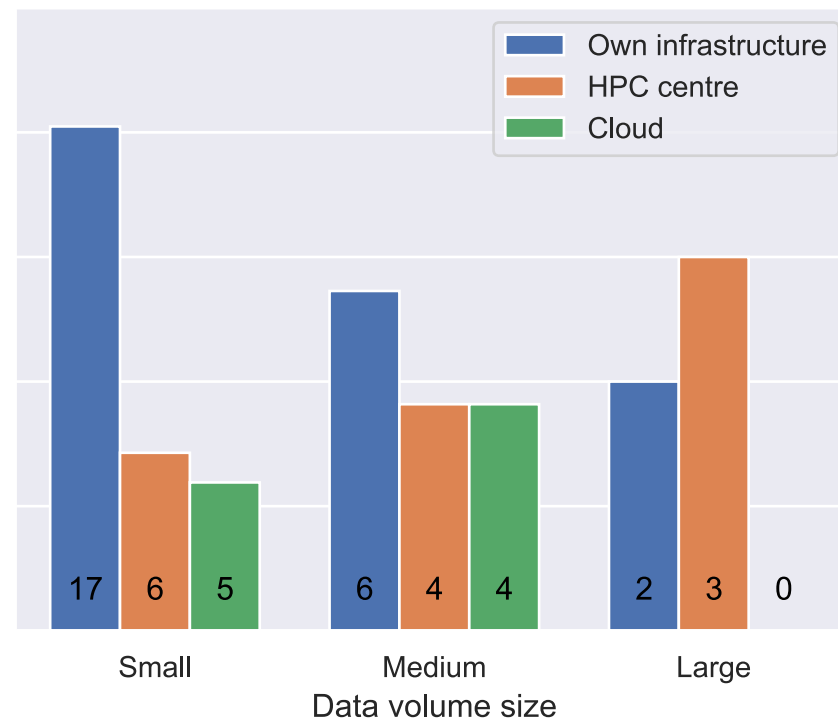
Data type to be processed



Programming Language to be used





Where to stage the data?




# Datacenter



(a) data

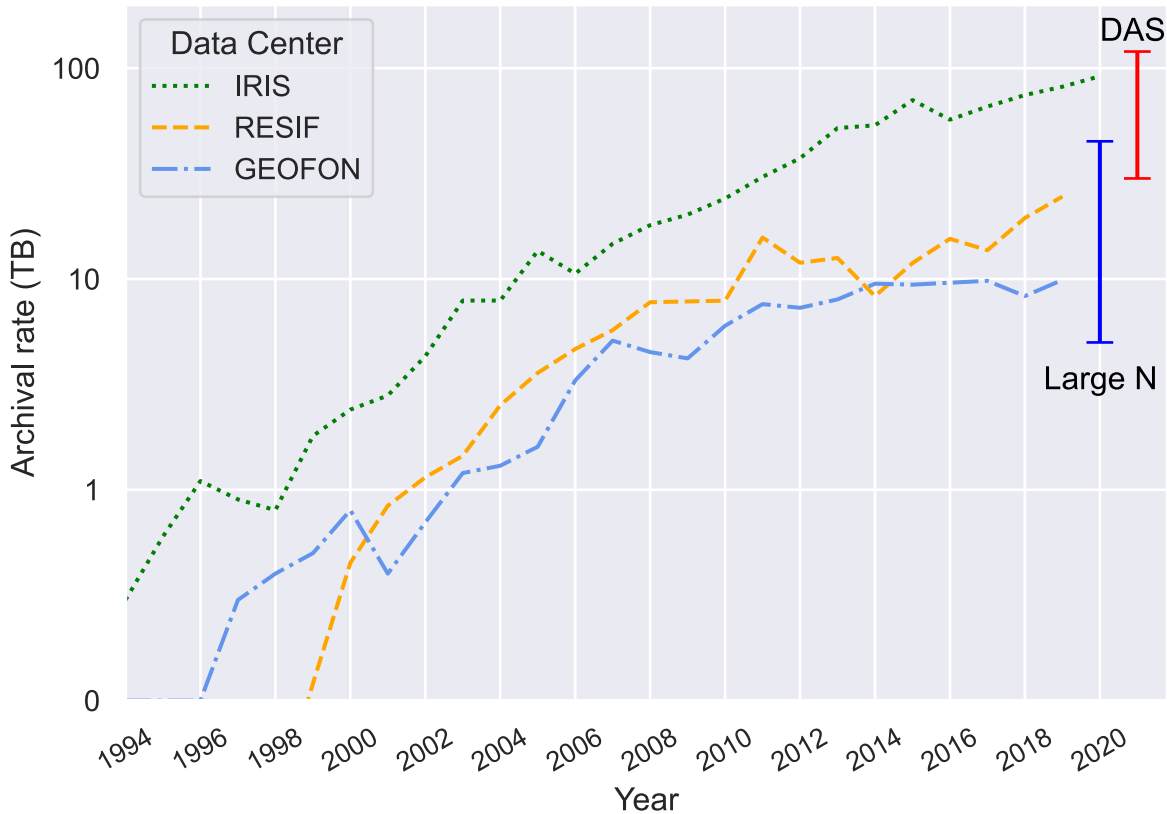


(c) data description



(e)

Archival rate per year



Percentage of requests by size

