International Environmental Modelling and Software Society (iEMSs)
7th Intl. Congress on Env. Modelling and Software, San Diego, CA, USA,
Daniel P. Ames, Nigel W.T. Quinn and Andrea E. Rizzoli (Eds.)
http://www.iemss.org/society/index.php/iemss-2014-proceedings

Technology behind the Deltares Open Archive

Grijze, Andre, Peter Gijsbers[1], Erik de Rooij[1], E. Pelgrim[1], Onno Van Den Akker[1]

[1] *Deltares, P.O.Box 177, 2600 MH Delft, The Netherlands (andre.grijze@deltares.nl)*

Abstract: Water agencies conducting daily operational forecasting need an archive solution that enables them to review (and defend) forecasts and decisions made, conduct post-event performance analysis, create canned datasets for training and create datasets for model calibration and hindcasting. This archive solution should store datasets that range from time series (observations and forecast), workflow definition and model-run settings, to web-reports and text products (e.g. forecaster notes). The challenge is to provide an open solution that supports both operations and science, allowing access by a variety of applications. The Deltares Open Archive is a new solution which meets those needs. The solution consists of a data storage, a catalogue and services to store, backup, discover and retrieve data.
The data storage is file based using standardized data formats such as NetCDF-CF for scientific data and XML, HTML, PDF, txt and zip to store non-scientific data. Associated metadata is held in xml. Data access is provided by the THREDDS Data Server (TDS), a web server that provides a variety of standardized remote data access protocols like OPeNDAP, WMS and WCS for scientific datasets and plain HTTP for non-scientific data. To search the Deltares Open Archive, a key-word based catalogue solution is developed using GeoNetwork, an OGC Catalogue Service compliant product. Metadata is composed of area, data source and temporal information, while more tags can be added to highlight e.g. threshold crossings and label a period as being a relevant (storm) event. Data can be provided by Delft-FEWS or any other application, as long as metadata requirements are met.

Keywords*:* Open; archive; scientific data; THREDDS, GeoNetwork, CSW

1. INTRODUCTION

Many water management agencies have operational duties that require data collection, data quality control and real time forecasting. Since 2003 Deltares has been working with over 50 agencies around the world to jointly develop real time information and forecasting applications based on the Delft-FEWS software platform [Werner et al., 2013; Deltares, 2014]. During these contacts water agencies expressed the desire for a data archiving solution that stretches beyond the capabilities of many other hydrological data archive solutions such as WISKI [Kisters 2014], Aquarius [Aquatic Informatics 2014] or HIS [CUAHSI 2014]. Water agencies conducting daily operational forecasting need an archive solution that enables them to (i) review (and defend) forecasts and decisions made during events, (ii) conduct post-event performance analysis, (iii) create canned datasets for training and (iv) create datasets for model calibration and hindcasting. This archive solution should store datasets that range from time series (observations and forecast), workflow definition and model-run settings, to forecast products and communication records (e.g. forecaster notes, email conversations) and event review reports. Dependent on the use case differences exist in the data to be archived and their lifetime. An in depth analysis of the use cases is discussed in Gijsbers et al. [2014].

To meet the requirements of these use cases Deltares developed the Deltares Open Archive. The Deltares Open Archive is an open, scalable archive based on OGC standards for searching, storing and accessing the data. This archive is capable of storing simulations, observations, configurations, messages, reports and rating curves. The concept of events is introduced to improve data discovery as well as life time management of data. The first application of the Deltares Open Archive is currently being deployed in a pre-production environment for the Australian Bureau of Meteorology. The Australian archive will initially be populated by migration of their legacy archive. In future their new Delft-FEWS based Hydrological Forecasting System will continue population of the archive. The second application may well be the follow-up of the Matroos forecast archive from the

Dutch Rijkswaterstaat. Currently an assessment is on-going for the potential migration of Matroos to the new Deltares archive solution. Other water management organizations also have shown interest and plan on using this archive solution.

This paper describes the technology behind the Deltares Open Archive, starting with a short analysis of the design principles applied and a high level overview of the architecture. Its major components (data store, data access and catalogue) will be presented in more detail followed by a discussion of the advantages and disadvantages of the solution chosen. The paper concludes with a summary and future work ahead.

2. ARCHITECTURE PHILOSOPHY, REQUIREMENTS AND DESIGN

The technical architecture of the Deltares Open Archive is inspired by the Open Earth philosophy as introduced by Koningsveld et al. [2010]. First of all it intends to be an *open* solution which means that the archive should be capable of storing time series and other data types from any source regardless of its data format and should be accessible by a wide variety of applications. To achieve openness and accessibility by a wide variety of applications, a solution based on industry based *standards* is required. The solution needs to facilitate different accessibility methods to accommodate scientists - preferring raw scientific data-, review teams - requiring all governance details of the products generated - and decision makers – requiring end products such as images or reports. Since an archive will grow over time, scalability is important to keep performance of both discovery and retrieval acceptable for the use case(s) at stake. Users and applications using the archive should be capable of finding the data with ease. Operational water agencies have indicated that retrieval of data associated with an event, e.g. a particular hurricane, is important in various use cases. Finally, water agencies have made clear that they have their own policies regarding hardware and operation systems to host the various applications within the organization. Platform independence is critical for success, while the technology stack chosen should have a proven track record including availability of support services before being accepted by the ICT-departments. These departments also expressed a desire for understandable solutions with limited complexity that requires no special skills to administer and maintain.

As shown in Figure 1, the technical architecture follows a layered design, separating storage from access, both for the actual data archived as well as the catalogue to find the data. The data storage is file based. It can store different types of datasets like simulations, observations, configuration/model setup, messages and rating curves. Each dataset consists of several files stored in its own directory and is accompanied with a metadata file in xml to describe its contents. While the time series are always stored in NetCDF files, other data types can be stored in a proprietary format. Data access is provided via remote data access protocols for file transfer (HTTP) and geospatial services (OPeNDAP, OGC WMS and/or OGC WCS). To find the data, a catalogue is populated which can be accessed by the OGC CSW protocol. The catalogue caches the metadata information in a geospatial database to facilitate quick access. A harvester transfers metadata files from the data store to the catalogue database. To accommodate time series data that are not available in NetCDF format, a data importer is available to read proprietary time series formats and convert the data to the NetCDF format for storage in the archive.

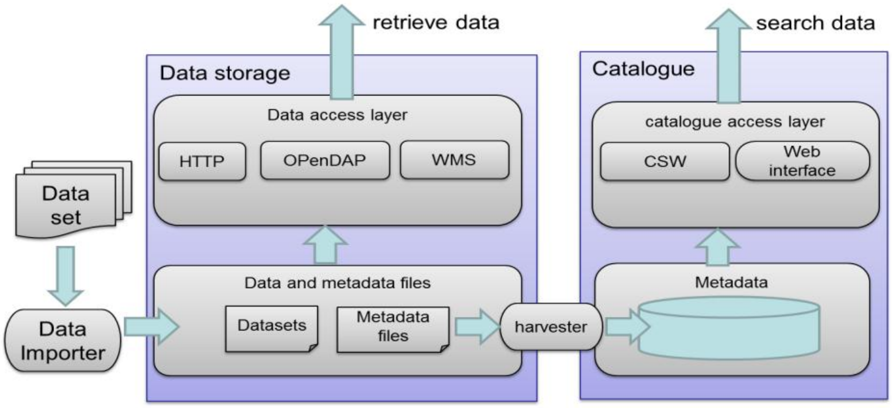

Figure 1 Architectural design of the Deltares Open Archive

3. DATA STORAGE

Data storage is a primary function of an archive. As indicated the data is stored in files directly on a disk in its own directory without using database software. The root of each directory contains a metadata file which describes the content of the dataset. Figure 2 shows the structure of the file based data storage.

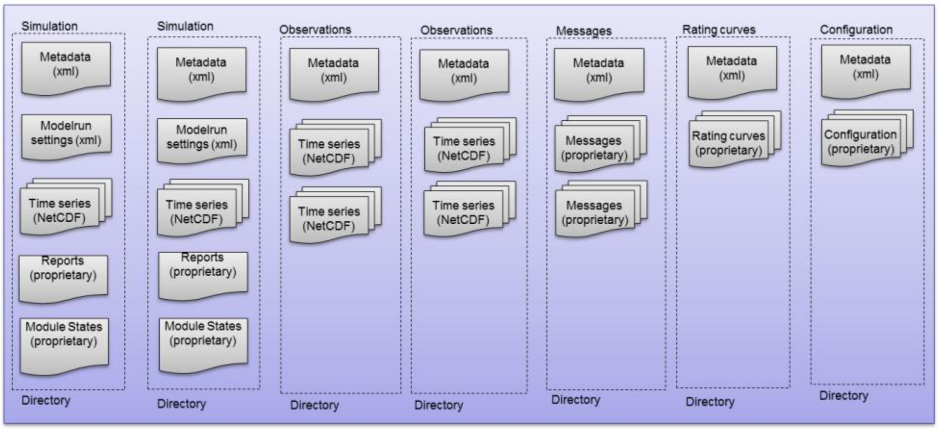

Figure 2 Structure of the file based storage

The Deltares Open Archive supports several kinds of datasets where each kind of dataset consists of one or more data types. Table 1 provides an overview of the supported datasets and the related data types.

Table 1 Overview of the supported datasets and data types

Data type / **Data set**	Time series	Model run settings	Model States	Reports	Messages	Rating curves	Configuration
Simulation	X	X	X	X			
Observations	X						
Messages					X		
Rating Curves						X	
Configuration							X

Based on system administrator preferences a self-explaining directory structure is generated using the date stamp, the area or source of the data and the kind of dataset. For example the observations for the Brisbane basin of 24 January 2014 will be stored in the directory <dataroot>/2014/01/brisbane/24/observations. This division by month and area of interest will enable system administrators to easily move data around to deep storage devices once data usage has become limited.

To meet requirements of openness, some data types will be stored in a prescribed data format (time series, metadata, run info) while others are allowed to be stored in proprietary formats. Table 2 provides an overview of the supported data formats for each data type. The metadata files and the model run settings are stored in a xml-file following a Deltares Archive Xml Schema Definition. Time series are stored in NetCDF files following the Climate Forecasting convention for definition of its temporal, spatial and variable dimensions. NetCDF is a standard data format designed for storing time series, both scalars and coverages, in an efficient, scalable, portable and platform independent way [OGC, 2010]. A wide variety of free and commercially licenced tools and libraries are available for accessing the content of a NetCDF-file. In addition, several open source server implementations provide standardized remote access to NetCDF files. See [Unidata, 2014a] for a complete software overview.

Table 2: Overview of the data types and the supported data formats

Data format / **Data type**	NetCDF-CF	XML	Proprietary format
Time Series	X		
Metadata		X	
Model run settings		X	
Model states			X
Reports		X	X
Messages		X	X
Rating Curves		X	
Configurations			X

Not all organizations have time series data available in the NetCDF-CF format. While an open archive should be capable of storing time series from any source regardless of its data format, a standards based approach provides advantages in terms of accessibility. To enable storage in NetCDF-CF, while accommodating other data formats, a Data Import utility is available as part of the Deltares Open Archive. This utility, based on Delft-FEWS code, can read about 200 different time series file formats and export these time series to the archive in NetCDF-CF. The utility also creates the required metadata files at the same time. When the archive is populated by a real time system, it is preferred to directly write the data to the archive data storage in the appropriate data formats.

4. DATA ACCESS

The data access layer provides remote access to the datasets. Non-scientific data can be downloaded by file using HTTP Get calls. Time series can be downloaded as file using HTTP Get or by a scientific data access protocol such as OPeNDAP, OGC WMS or OGC WCS. All services are implemented using the THREDDS Data Server. THREDDS is a web service implementation, developed and supported by Unidata, which can be deployed in any java-application server [Unidata, 2014b].

OPeNDAP provides access to subsets of time series data and can also aggregate time series data over several NetCDF-CF files. OPeNDAP is a data transport architecture and HTTP based protocol widely used by scientists [OPeNDAP, 2014]. Data hosted by OPeNDAP can be stored in different locations (e.g. over different files) hidden to the clients. In addition, it is capable of hosting large collections of data, and it is a proven, widely used and accepted solution supported by various open source tools. The OGC Web Coverage Service [OGC, 2009] is another protocol available to retrieve raw geospatial time series data. When requesting the data using the OGC Web Map Service protocol [OGC, 2006] a geo-referenced map image (e.g. JPG) is generated from the geo-spatial time series data and handed over to the client.

5. DATA DISCOVERY

While data storage is a primary function of the archive, its value increases significantly when people can find the data they need. Metadata, describing the content of the dataset, is crucial. Based on this metadata a catalogue service can be built to facilitate data discovery, turning the data storage into a valuable asset. Each dataset in the Deltares Open Archive data storage is described by a metadata xml file following a straightforward Xml Schema Definition. Primary metadata items to describe a dataset in the Deltares Open Archive catalogue are (i) kind of dataset (ii) start date/end date of period covered (iii) basin or area of interest (iv) locations and variables of available time series (v) source of data/production method (vi) organisation that generated data (vii) a list of URLs to files in the dataset.

While the above metadata items provide the basics to describe the data and find data time series for general modelling and data purposes, it does not facilitate some of the discovery tactics that water agency staff would like to follow to support their use cases [Gijsbers et al, 2014]. Water agencies are interested to find out where thresholds are crossed in a particular area, when thresholds are crossed at a particular location or what data is available associated with a certain (flood) event. Two capabilities have been added to the Deltares Open Archive solution to facilitate those needs. First of all an open set of key-value pairs can be added to the metadata. An observations dataset that includes a period of flooding at e.g. Amberley Station can then be identified by extending a metadata record with a set of key-value pairs (e.g. location=Amberley Station, variable=Observed water level, threshold=minor flood). Secondly, the concept of events is introduced where an event is defined by a name and description, a start and end time and an area where the event takes place. Examples can range from flood events or storms to spill calamities or prevailing weather conditions (droughts or winds). All datasets that fit within this period and area are considered part of the event dataset. Events can thus be seen as a container concept for multiple datasets. Event tagging is done via a GUI as the data user can assess whether a period of record is of sufficient interest for a particular use case.

Within the Deltares Open Archive, the catalogue service is built using GeoNetwork. GeoNetwork is an open source solution which is capable of storing and hosting metadata records for datasets according to several industry standards for metadata and catalogue services [GeoNetwork, 2014].The metadata records are stored in a geospatial database to enable catalogue services based on e.g. the OGC standard Catalogue Service for the Web (CSW) [OGC, 2007]. CSW is based on the HTTP-protocol and can therefore be accessed from any programming language. CSW supports spatial and keyword based queries allowing the kind of queries that are needed by the water agencies. The first client application available to query the catalogue is fully keyword based and embedded into the Delft-FEWS operational forecasting platform. While not yet implemented, the GeoNetwork technology stack allows extension of the search functionality into the spatial domain by populating the underlying geospatial database with the required spatial information.

As indicated, each dataset is accompanied by a metadata xml file following a schema of the Deltares Open Archive. The catalogue is populated from the information in these metadata files by a harvester. The harvester reads the Deltares metadata file and creates a metadata record according to the ISO 19139 standard in GeoNetwork catalogue for each dataset found in the data storage. The harvester runs periodically, typically every night. Each time, the harvester synchronizes the catalogue based on the latest status of the datasets and metadata files in the data storage. The catalogue holds a record-id for each dataset in the catalogue. This record-id is also stored in a file in the directory of each dataset. This ensures that the harvester can properly transfer metadata updates at the data store to the catalogue. When the harvester scans the datasets the following situations can occur: (i) a dataset without a record-id is found. In this case a new metadata record will be created; (ii) a dataset has a record-id while the metadata file is updated after the last scan of this dataset. In this case the harvester updates the metadata record in the catalogue using the record-id as a primary key; (iii) a dataset has a record-id while the metadata file is not updated after the last scan of this dataset. In this case nothing is done because the catalogue is already up-to-date; (iv) a record-id without a dataset is found. In this case the record will be deleted from the catalogue.

6. ARCHIVE MANAGEMENT

System administrators need to preserve the integrity and accessibility of the archive while keeping its volume manageable given hardware capabilities. Within the design of the Deltares Open Archive, it is assumed that integrity of the data store (datasets and associated metadata.xml files) can be preserved using normal file backup procedures. When any of the files in the data store is corrupt, the backup can be restored. The recovery strategy for a corrupt catalogue database can be conducted in three steps (i) drop the database instance, (ii) remove all record-id files from the data store and (iii) rerun the harvester to populate the catalogue with new metadata records.

The volume management strategy is related to the use cases of the data stored as different pieces of data can have different lifetime requirements. Generally, data related to an event is considered more important than data related to a ‘business-as-usual’ situation. Data not related to an event should be removed or archived to tape after a certain amount of time. The Deltares Open Archive comes with a data management tool which reports, based on configurable set of rules, what data can be removed or moved to tape. System administrators and data users together have to define the set of rules, namely the default lifetime of a dataset and the lifetime of an event. When a dataset is labelled as part of an event, the lifetime of the event is assigned to the dataset as the event lifetime overrules the default lifetime of the dataset.

7. DISCUSSION OF IMPORTANT DESIGN CHOICES

The Deltares Open Archive has been developed after consultation with its target audience: hydrologists and system administrators at water agencies with operational duties involving data collection, model development, forecasting and associated decision making. Their preferences for a solution that is easy to use and maintain have played an important role in design choices made. This section highlights a few of the advantages and disadvantages of the design chosen.

The Deltares Open Archive uses a file based storage systems as these are simple, flexible and scalable. If well-organized they are also easy to manage. Alternatives could be a relational database management solution or a *big data* solution. Although relational database management systems (RDBMS) can cope with time series by storing one row per time stamped data entry, performance is generally poor and storage is inefficient, especially in the case of grid data. The only way to keep RDBMS performance sufficient when storing time series is to store a large amount of data entries in a single record, for example by storing the values in a Binary Large Object (BLOB) [Deri et al., 2012]. However, this approach will only allow access to the time series by using a specialized tool to convert the BLOB to a time series which could, dependent on availability, break the requirement for an open archive. While hydrological archives can be massive in volume, they are typically well structured. At the time of writing of this paper several big data-solutions, like OpenTSDB or Cassandra, are available which are capable of storing large amounts of (un)structured data including time series [OpenTSDB, 2014] [Cassandra, 2014]. These solutions can provide good performance even when storing huge amounts of data. This solution direction has not been chosen since the technology is relative new and the knowledge required to host such solutions is not available in the ICT departments of many water agencies. Another disadvantage is the limited 'openness' of a big data database with little open source tools adopted by the water community. Although most big data solutions offer the possibility to query the data by using a SQL-like language, the range of data access options is considered limited compared to the geospatial capabilities offered by solutions such as THREDDS.

The Deltares Open Archive uses a keyword based search strategy to discover interesting datasets. Queries based on the data values in the dataset itself are not accommodated, partly because GeoNetwork is not very well suited for these types of queries. This can be considered a disadvantage, as searching for absolute value crossings often comes to mind. SQL-based archive solutions often allow such queries. From a hydrological data search perspective however, a search based on absolute values is often only done by local experts as you need to know the local situation to know at what level a location is flooding. If a data user not having this local expertise wants to retrieve observed data from the archive for all locations which were in flood during a certain period of time, it becomes much easier if he/she can search for all datasets which are tagged with the keyword ‘threshold=major flood’. Such metadata tag can be added during data production based on general rules as defined by the local export. Alternatively, a data crawler may be developed that uses similar rules to enrich metadata at some later stage.

Within the Deltares Open Archive the point of truth for metadata is the metadata xml file that accompanies each dataset. The catalogue is built from these xml files. It seems odd to store the metadata both in the catalogue and with the dataset itself, but this approach has several important advantages. First of all, the approach allows the catalogue to be built from scratch in case of a corruption of the catalogue. This enables recovery of the catalogue after any operational problem. Secondly, this approach enables switching to another catalogue solution without affecting the data archive itself.

8. SUMMARY AND FUTURE WORK

This paper presents the Deltares Open Archive, an open archive solution for water management agencies dealing with operational data collection and forecasting tasks. The solution offers a file based data storage mechanism which allows easy management of growing data volumes by storing important (e.g. event related) data on disk while moving other data to tape over time. To find data in the archive, a catalogue web-service is provided with a focus on keywords to query for the period and area of interest, thresholds being crossed and/or its relation to a particular water related event. These concepts fit closely to the specific use case requirements for real-time event comparison, post-event reviews and inquiries, training. Alternatively, geospatial queries could be composed when associated spatial data is made available to the catalogue. Data retrieval can vary from HTTP based file downloads to scientific data access using geospatial web-services.

Future work includes the development of a web interface for system administration purposes to provide system status overview and allow system administration processes to be executed on demand. Other activities foreseen include a shared web-service for searching and data delivery, as well as extensions to the harvester and the data importer. Existing datasets often are poorly searchable due to lack of search keys. The extension foreseen is a plugin-mechanism for the data importer and the harvester that allows developers to write their own custom code for analysing datasets and creating extra metadata tags.

9. REFERENCES

Aquatic Informatics 2014. Water Data Management Solutions http://aquaticinformatics.com/products (last accessed 12.03.2014)

Cassandra 2014. Cassandra http://cassandra.apache.org/ (last accessed 11.05.2014)

CUAHSI 2014. CUAHSI HIS sharing hydrologic data http://cuahsi.org/his.aspx (last accessed 12.03.2014)

Deltares 2014. Delft-FEWS Community Portal. http://www.delft-fews.com (last accessed 12.03.2014)

Deri, L., Simone, M., Francesco, F., 2012. Tsdb: A Compressed Database for Time Series. Download available at http://luca.ntop.org/tsdb.pdf

Gijsbers, P., Grijze, A., Dijk, M. van, Akker, O. van den, 2014. Usage of the Deltares Open Archive. In: Ames, D.M. and Quinn, N. (Eds.) Proc. of the 2014 Int. Congress on Environ. Model. and Software (iEMSs), San Diego, California, USA Session A2: Sharing Scientific Environmental Data and Models.

GeoNetwork 2014 GeoNetwork opensource. http://geonetwork-opensource.org/ (last accessed 12.03.14)

Kisters 2014. WISKI Your Hydrological workbench http://www.kisters.eu/water/software.html (last accessed 12.03.14)

Koningsveld, M. van, Boer, G.J. de, Baart,F., Damsma, T., Heijer, C. den, Geer, P. van, Sonneville, B. de, 2010. OPENEARTH – Inter-company management of data, models, tools & knowledge. WODCON XIX, Beijing, China. Download available at: https://publicwiki.deltares.nl/download/attachments/042401895/VanKoningsveld+et+al.+(WODCON2010)+Data++Models+and+Tools+-+V17.pdf (last accessed 12.03.2014)

OGC 2009. Web Coverage Service. http://www.opengeospatial.org/standards/wcs (last accessed 12.03.14)

OGC 2006. Web Map Service. http://www.opengeospatial.org/standards/wms (last accessed 12.03.14)

OGC 2007. Catalogue Service. http://www.opengeospatial.org/standards/cat (last accessed 12.03.14)

OGC 2010. OGC network Common Data Form (netCDF) standards suite. http://www.opengeospatial.org/standards/netcdf (last accessed 12.03.14)

OPeNDAP 2014. OPeNDAP. http://www.opendap.org/ (last accessed 12.03.14)

OpenTSDB 2014. OpenTSDB. http://opentsdb.net/ (last accessed 11.05.2014)

Unidata 2014a. Software for Manipulating or Displaying NetCDF Data. http://www.unidata.ucar.edu/software/netcdf/software.html (last accessed 12.03.2014)

Unidata 2014b. THREDDS Data Server (TDS). https://www.unidata.ucar.edu/software/thredds/current/tds/ (last accessed 12.03.14)

Werner M., Schellekens, J., Gijsbers, P., Dijk, M.van, Akker, O. v.d., Heynert, K., 2013. The Delft-FEWS flow forecasting system. Environ. Model Softw. 40 (2), 65-77