

Request for Information RFI 1501-18 (OBS) Cache in and through the cloud

« Cache in the cloud » team
January 2018



WMO OMM

World Meteorological Organization
Organisation météorologique mondiale

Purpose and Objective

- This RFI is a procurement exercise to conduct a market survey in order to obtain information (from the market) that can be used to identify available or potential solutions for fulfilling identified needs which may include information on **cost (budget)** and delivery time;
- The main purpose of this RFI is to request vendors to provide an estimated not-to-exceed (NTE) cost for this project and is intended only for budgetary purposes;
- This RFI is entirely non-binding to the WMO and does not represent a commitment of any kind towards any vendor who will respond to this RFI;
- Considering WMO context, it is important for the interested members to know in advance the NTE amount of the project;
- This presentation describes the project from a high level perspective;
- The result of this RFI will then be used to decide whether it is commercially and technically viable to go for an official tender;
- WMO (World Meteorological Organisation) plans to launch an official tender for the benefit of its members in late 2018;
- This RFI does not constitute a commitment from the vendors that their proposal, if they decide to bid, will be capped by this amount;



What is the WIS ?



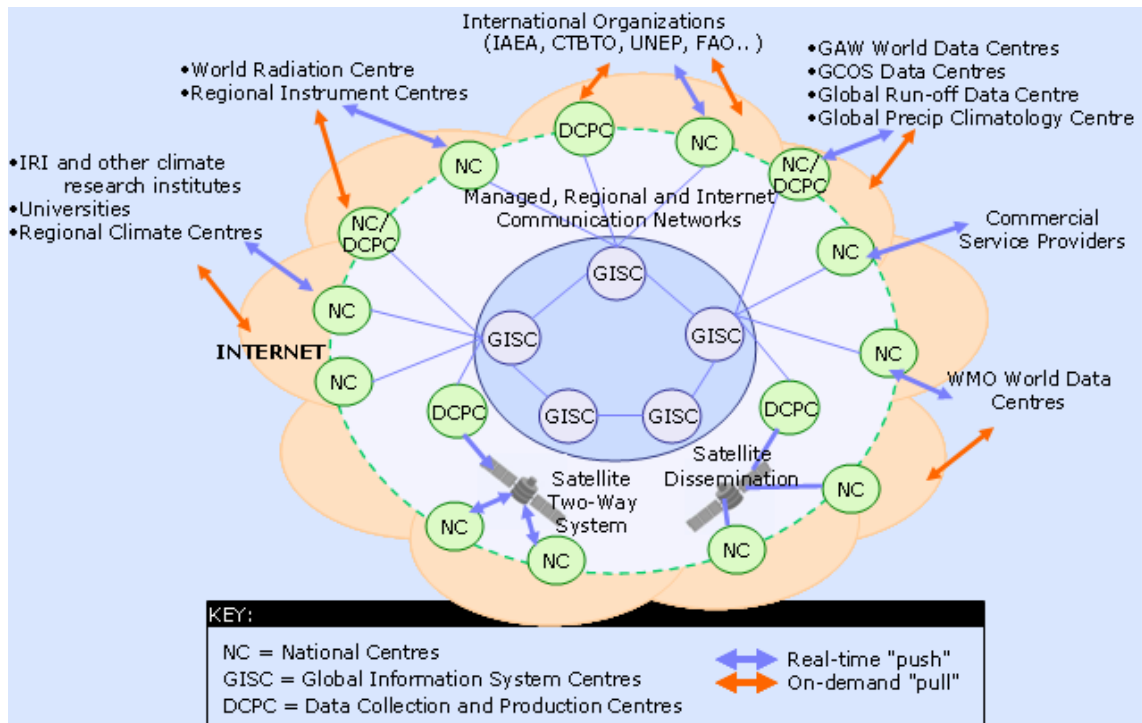
WMO OMM

A brief description of the WIS

- The WMO Information System (WIS) is the name of the service used within the meteorological services around the world to exchange and then make available data and metadata to their users
- It is based on 15 centres (called GISC) around the world that have a particular responsibility in the system.
- The GISC must collect data from centres in their area of responsibility and distribute the data with global interest to all other GISC to constitute a cache of data
- The global cache then exists in 15 versions (one in each GISC)
- The GISC support the connected centres in the creation and maintenance of the metadata records.
- Metadata is then exchanged between all GISCs to maintain an identical catalogue at each GISC
- Further information can be found on WMO website:
<http://www.wmo.int/pages/prog/www/WIS/overview.html>



A simplified view of the architecture



List of the GISCs

- Beijing (China)
- Brasilia (Brazil)
- Casablanca (Morocco)
- Exeter (UK)
- Jeddah (Saudi-Arabia)
- Melbourne (Australia)
- Moscow (Russia)
- New Delhi (India)
- Offenbach (Germany)
- Pretoria (South-Africa)
- Seoul (Republic of Korea)
- Tehran (Iran, Islamic Republic of)
- Tokyo (Japan)
- Toulouse (France)
- Washington (USA)



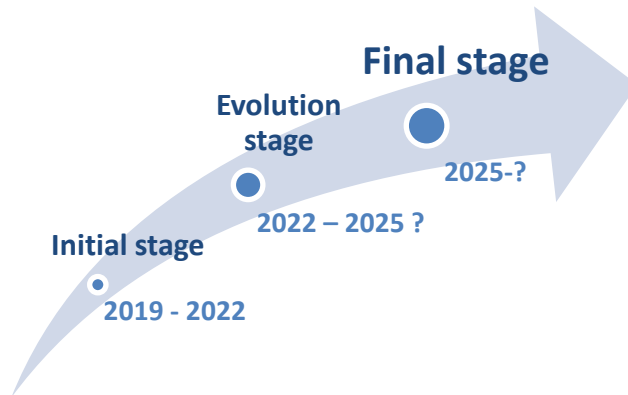
GISC, cache and metadata within WIS

- Among other activities, the GISCs have the responsibility to maintain a cache of data for the last 24 hours and the catalogue of all data available within WIS
- This RFQ is about a service that will facilitate the GISC to maintain both the cache of data and the metadata catalogue
- The only users of the service will be the GISCs.
- The required protocols mentioned in this document are the protocols currently used by the GISCs to fulfil the functional requirements of the WIS.

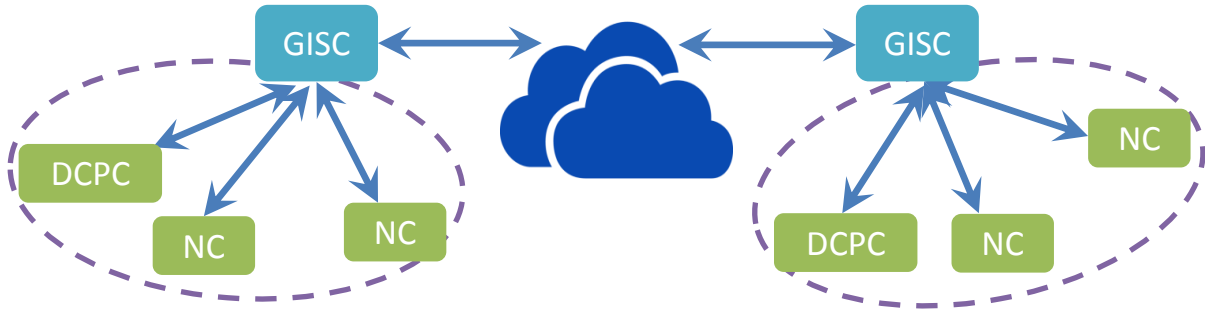


A step toward WIS 2.0

- The WIS is currently evolving toward a new architecture where in addition of the data exchange, services will be proposed by the WIS in its version 2.0
- It is considered that the cache in the cloud is one of the first step in the implementation of the WIS 2.0 system
- Considering the timeline of the implementation of WIS 2.0, the contract for the desired service will last 3 years 2019 - 2022



Shared cache between GISCs



- Each GISC collects data from its connected centres (the collection of data from these centres is out of scope of the project)
- The data intended for GlobalExchange must then be available in a cache kept locally by each GISC.
- The service required will facilitate the exchange of data between the GISCs to create the 15 local caches



Metadata catalogue

- Each GISC helps the data producers within its area of responsibility to describe their data with metadata (ISO19115 compliant) record
- The GISC harvest the metadata catalogue from other GISCs using OAI-PMH and establish a global catalogue



Project timeline

- The results of this RFI will be discussed with various WMO bodies in the first half of 2018.
- If the project is approved, it is anticipated that the official tender will be launched in Q4-2018
- The service will then start for a 3-year period in Q2-2019
- The tender will only cover the « Initial Stage »



Summary: what are we looking for?

A fully managed cloud based service to support the GISCs in creating the shared cache and the metadata catalogue as described in the WMO Information System



Functional and technical requirements



WMO OMM

Guiding principles

- Adherence to WIS technical specifications (maintain technical interfaces currently in places - architecture, protocols,...).
- Evolutionary – support current protocols, evolve to support new ones over time, seek opportunities for improving efficiency
- Address both data and metadata exchange requirements
- Consider governance, performance, reliability, security, administration, monitoring, and financial aspects of the project
- Must respect local legislation of each GISC in term of networking connectivity



Data requirements in the cloud

- Performance
 - Global data is to be available at all GISCs within 15 minutes. The largest files to distribute are less than 200MB.
 - High priority data is to be available at all GISCs in 1 minute. The data in this case is very small (a few kilobytes maximum).
- Scalability
 - The total volume cache is estimated to be less than 200GB per day in files varying a lot in sizes (from less than 1kB to 200MB).
 - The number of files per day is above 100,000.
 - Each GISC contributes to a fraction of the volume. However, the volume provided by each GISC varies greatly (a factor of at least 20 exists between heavy contributors and some others)
- Data exchange
 - Must use agreed protocol to exchange data: FTP, SFTP, HTTP, HTTPS for GISCs
- Persistence of the cache
 - In addition to receiving and pushing the data in real time to the GISC, a copy of the data will be made available so that GISC can re-access old (up to 5 days old) individual files received by all GISCs (using the agreed protocols FTP, SFTP, HTTP, HTTPS). In total this cache will be less than 1TB (200GB per day x 5 days)



Metadata requirements in the cloud

- Metadata exchange
 - Must use agreed protocols to exchange metadata (OAI-PMH)
 - Will harvest once a day the updates of catalogue of all 15 GISCs
 - Will make this catalogue available for harvesting by all GISCs
- Once a day, a compressed copy of the catalogue will be prepared and made available to the GISCs using FTP, SFTP, HTTP or HTTPS
 - The GISCs will be able to download this compressed copy of the catalogue
 - The volume of the compressed catalogue will be less than 1GB



Administration requirements

- Configuration and management of IT environment
- Configuration and management of software packages that are required by the service
- Management of access rights
 - A very small number of systems and users will need to have access to the service
- Activity monitoring – The service should propose a mechanism (API, access to logs,...) so that GISCs can integrate the service monitoring into their own system
- Data availability and timeliness
 - Monitoring & statistics



Security requirements

- The security of the service is a must
- State of the art measures must be in place in order to guarantee that only the GISCs will have access to the service. This includes, but is not limited to:
 - Access control lists/firewall rules so that only known GISCs can access the cache service.
 - Only advertise the protocols needed to run the service
 - “Harden” the platforms so they do not run unnecessary services
 - Patching and updating policies
 - ...
- The design should also include anti-DDOS service



Networking requirement - RMDCN

- All GISCs except Brasilia, Tehran and Washington are connected to a private MPLS network, provided by Interoute, called the RMDCN (<https://www.ecmwf.int/en/computing/our-facilities/rmdcn>)
- The RMDCN has a very high SLA
- It is anticipated that many GISCs may prefer to use the RMDCN for the exchange of data
- Ideally, an interface between the cloud service and the RMDCN should be proposed. Vendors are invited to confirm whether they can connect the proposed solution to the RMDCN and to include this in the « not to exceed » figure.



Networking requirements

- The Internet will be used for the exchange of data and metadata between the cloud service and the GISCs, either exclusively or as a complement to the RMDCN
- Each GISCs will have the responsibility to get its own (normally redundant) Internet connection
- The data centres hosting the service must have redundant Internet access
- All cloud servers and IT environment required by the service must be interconnected by a highly reliable networking solution. The anticipated solution must be described.



Required protocols

- The protocols that can be used between the GISCs and the cloud service are:
 - FTP, SFTP, HTTP, HTTPS for data
 - OAI-PMH for metadata
- In addition, it is possible to use additional protocols within the cloud service to distribute the data and metadata between the various components of the service



SLA

- The WIS is a critical system that is used 24/7 by all National Meteorological Services around the world.
- Therefore, the « Cache in the cloud » service must have a very high SLA – 99.95% (including planned maintenance)
- It means that the service must be designed in a highly redundant manner in order to achieve the required SLA



A potential architecture



WMO OMM

Information

- The following slides present a potential architecture and software solution that can be used in providing the service
- They are based on practical experience from the members
- However, it is the responsibility of the provider to design the solution in order to fulfil the requirements
- As such, these are only indicative and shouldn't be considered as the requested design



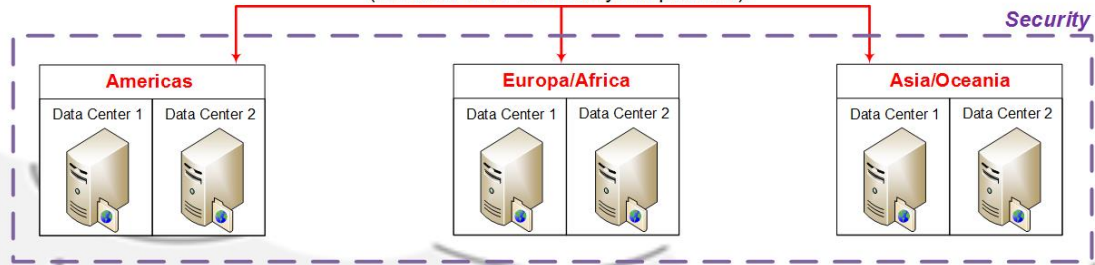
Example of IT architecture

- Considering the geographical repartition on the GISCs, a solution can be based on:
 - Two servers in two different data centres per “region”
 - Three regions:
 - Americas for the collection of data/metadata of GISC Brasilia and Washington
 - Europe/Africa for GISC Casablanca, Exeter, Moscow, Offenbach, Pretoria, Toulouse
 - Asia/Oceania for GISC Beijing, Jeddah, Melbourne, New Delhi, Seoul, Tehran, Tokyo
- Each GISC send data one server in its region. A failover mechanism is in place to fall back to the second server in case of an issue
- Each server in a region fetches metadata record from the GISCs in the region (once a day)
- Servers synchronize data and metadata within the cloud service
- One server send data to all GISCs within the region
- GISCs fetch metadata on a server of its region (once a day)

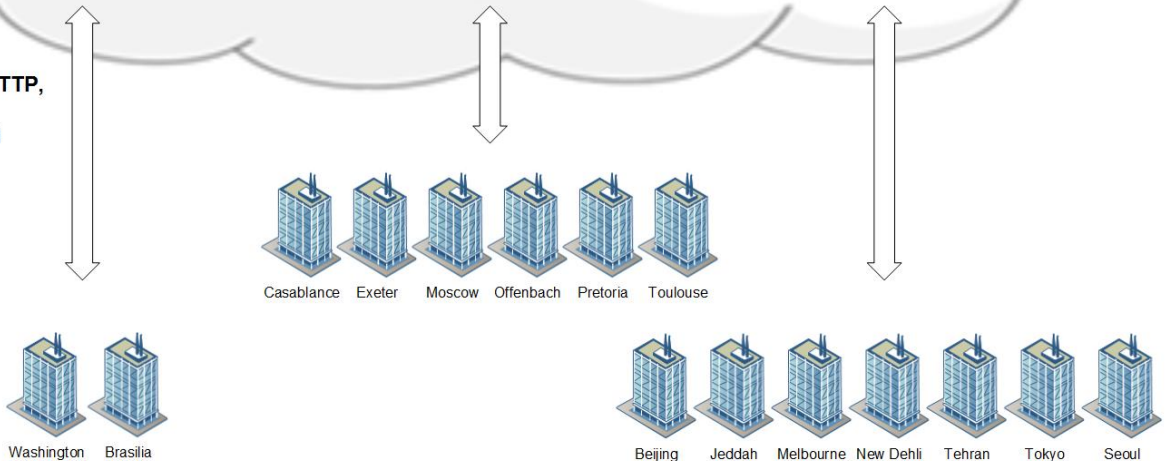


A possible design

Internal exchange within cloud = data and metadata
(Protocol to be defined by the provider)



Data = FTP, SFTP, HTTP, HTTPS
Metadata = OAI-PMH



Example of software

- AFD (Automatic File Distributor, <https://www.dwd.de/AFD/>) is an open source software widely used within the WMO community
 - It can be used to redistribute data using FTP, SFTP or get data using HTTP, HTTPS
 - It also handles to notion of urgent message
 - It is known to be very reliable
- pycswh (<http://pycsw.org/>) is an open source software to handle metadata records
 - It implements OAI-PMH
 - It is a very lightweight and simple to deploy software
- These are just examples of software that are known to be able to do the required functions. Using them is not a mandatory requirement

