

IGS Network & Data Center Position Paper

Darmstadt Workshop – May 2006

A. Moore, Jet Propulsion Laboratory, California Institute of Technology, CA, US

C. Bruyninx, Royal Observatory of Belgium

C. Noll, Crustal Dynamics Data Information System (CDDIS), NASA GSFC, MD, US

M. Scharber, Scripps Orbit and Permanent Array Center, CA, US

1. Introduction

We focus in this position paper on a specific plan for defining the IGS data flow in such a way to solve several actual problems that exist in the IGS infrastructure. This plan requires agreement and action from all stations' operational centers, all Data Centers, of the Central Bureau, and quite possibly of Analysis Centers. We have set up relatively simple rules that we feel will address actual problems in the IGS data flow without undue pain to participants and users.

Our motivation for writing a new specification of IGS data flow focused on the fact that customers of both hourly and daily IGS data expect that an outage of a single Data Center (DC) does not necessitate manual intervention to provide the data at alternate DCs. Therefore, we seek built-in redundancy with a secondary path of data flow at all IGS DC levels.

When a data file is later revised (which we allow for, but encourage engineering to minimize the occurrence), all DCs should end up with the proper latest version, again without manual intervention. It has become clear that in many cases so far, the version archived on the various DCs is not the same. Putting a unique version number on files is the most customary and correct way to avoid this confusion, because DCs can easily distinguish older versions from newer. However, in the beginning of the IGS, the need to include this feature was not obvious, due to the low number of stations and semi-manual analysis. IGS archival evolved without version numbers to the present day, when we see a large community depending on the IGS dataset in an automated fashion. We therefore do not consider a change to include version numbers here, due to its large impact on end users and hence, low probability of approval and adoption. Furthermore, informal conversations with Data Centers have suggested that some have limited resources to alter existing procedures, and it is not possible to consider algorithms that involve detailed comparisons of newly-arrived vs. already-arrived files. We also note that various conventions within the IGS, such as permitting DCs to choose not to archive all sites' data (within certain requirements¹), makes otherwise obvious recommendations such as using standard archive synchronization software likely to need unusual amounts of engineering in the IGS setting. Hence, we change the official IGS data flow to a PUSH-ONLY method (DCs currently both push and pull data to populate their archives). The proposed logic that DCs must follow is quite simple, and therefore has a high likelihood of being acceptable to DCs. If the rules outlined in this position paper are followed, the copies archived on the various DCs are guaranteed to be equivalent to those generated by the original station or its operational center, even if a data file is replaced much later than it was originally issued.

¹ See "Charter for IGS Data Centers" via <http://igscb.jpl.nasa.gov/organization/orgindex.html>

2. Definitions

Data : RINEX data pertaining to a single IGS station, possibly compressed using a combination of gzip, Unix compression and/or Hatanaka RINEX compression.

Data Centers (DCs): make data available to the public and include Global Data Centers (GDCs), Regional Data Centers (RDCs), and Operational Data Centers (ODCs).

Operational Center (OC): The agency that is responsible for making a station's data available to the IGS community. The details of getting the data from the station(s) may vary. An OC may or may not be an ODC, depending on whether it makes data publicly available. A station may be its own OC in some cases. An OC pushes data to DCs.

Operational Data Center (ODC): an OC that offers public access to data from a set of stations, usually those managed by itself or a partner agency.

Regional Data Center (RDC): collect and offer data from many stations and agencies in a region. The RDC must allow data to be pushed to it from OCs/ODCs. RDCs will push data to GDCs.

Global Data Center (GDC): collect and offer data from all IGS Reference Frame stations, as well as offering products (see DC charter). GDCs must allow data to be pushed to them from OCs/ODCs, RDCs, and other GDCs. GDCs will push data to other GDCs to equalize data holdings of IGS Reference Frame stations (at a minimum¹).

Upstream: the direction of flow of data (from station to final destination) as shown in Figure 1.




Figure 1: Flow of IGS data

Channel: Data flows between DCs on one of several channels. The sender of the data selects the proper channel according to flowcharts presented below, and transmits data to an upstream DC on that channel. In practice, channels are likely to be distinguished by different dropoff account names, or different incoming file subdirectories.

Replacement: The situation when an OC discovers that a file which has been transmitted upstream was deficient in some correctable way. The corrected file must replace the original file at all DCs. It is assumed that OCs set up automated quality control procedures that limit the necessity of replacements, so that replacements are the exception rather than the norm.

3. Setup

OCs/ODCs will associate with two DCs of any level: one primary and one secondary. The least upstream DC will be the primary. These DCs will be listed in section 13 of each station's IGS site log (see example in Figure 2 below).

Each RDC will define two GDCs as its primary and secondary upstream DCs. This will be coordinated by the DCWG.

```

13.  More Information

      Primary Data Center      : RDC1
      Secondary Data Center    : GDC2
      URL for More Information :
      Hardcopy on File
      Site Map                  :
      Site Diagram              :
      Horizon Mask              :
      Monument Description      :
      Site Pictures              :
      Additional Information    :
      Antenna Graphics with Dimensions

```

Figure 2: Example section 13 from IGS site log

Example:

OC	primary DC	secondary DC
-----	-----	-----
AAAA	RDC1	GDC1
BBBB	RDC1	RDC2
OC1	RDC2	GDC2
OC2	ODC2	RDC2
DC	primary upstream DC	secondary upstream DC
--	-----	-----
RDC1	GDC1	GDC2
RDC2	GDC1	GDC3
ODC1	RDC1	RDC2
ODC2	RDC1	GDC1

For this example, AAAA and BBBB are stations that act as their own OCs.

Each DC must be able to automatically determine whether it is a primary recipient for a given data file submitted to it by setting up, with downstream providers, data drop-off rules that identify the pushed data are on a principal, backup, or replacement path. GDCs must also be able to distinguish data submitted from other GDCs vs. data submitted from downstream sources. For example, a GDC may provide separate dropoff accounts such as

igsprincipal@gdc1.abc.xyz
igsbackup@gdc1.abc.xyz
igsreplacement@gdc1.abc.xyz
igsgdcmirror@gdc1.abc.xyz

Alternatively, a DC could use a single account, but different dropoff subdirectories:

<ftp://gdc1.abc.xyz/pub/incoming/igsprincipal>
<ftp://gdc1.abc.xyz/pub/incoming/igsbackup>
<ftp://gdc1.abc.xyz/pub/incoming/igsreplacement>
<ftp://gdc1.abc.xyz/pub/incoming/igsgdcmirror>

The GDC can clearly detect whether a pushed data file is on the principal, backup, replacement, or gdcmirror channel according to which area the sending DC placed it in.

Note that the use of protocols other than ftp transfer is not prohibited. DCs may, in principle, mutually agree to perform file transfer between one another by any means.

4. Rules

The following flowcharts illustrate the logic of these data flow concepts for OCs, ODCs/RDCs and GDCs.




Figure 3: How to be an OC




Figure 4: How to be an IGS ODC/RDC: Flowchart for when data for a site the ODC/RDC archives is received.




Figure 5: How to be an IGS GDC: Flowchart for when data for a site the GDC archives is received.

OCs and **ODCs** will always push station data to their two DCs simultaneously or in immediate succession. The principal channel will be used to push the data to the primary DC and the backup channel will be used to push the data to the secondary DC. If a data file is later corrected, the OC will transmit the corrected file to both DCs on the replacement channel.

All DCs must sleep (wait) 5 minutes before pushing upstream a file received on a backup path. This rule is presumed in all rules below as well.

RDCs will always push all IGS data that they received from downstream to their two upstream DCs. If received on the principal channel, the data is archived and pushed upstream to its primary DC on the principal channel and to its secondary DC on the backup channel. If received on the backup channel, the same occurs after a 5 minute wait, if a copy has not already been received on the primary channel. If a file already received is received again, it is discarded unless received on the replacement channel. In that case, the previous version is overwritten and pushed upstream to both DCs on the replacement channel.

GDCs will archive and push all IGS data received on a principal channel to all other GDCs. A GDC will also archive and push data received on the backup channel to other GDCs after a 5 minute wait, if a copy has not already been received. If a file already received is received again from a downstream station/DC on the principal or backup channel, it will be discarded. If a file already received is received again on the replacement, it will overwrite the previous version and be sent to all other GDCs on the gdcmirror channel. A file re-received on the gdcmirror channel overwrites previous versions and is finally archived by the GDC. GDCs may discard data from stations they do not archive, recalling that GDCs must minimally archive all IGS Reference Frame stations¹.

In addition, any IGS station data file “pulled” by a DC becomes a private copy and must not be offered on public IGS DC areas or pushed to upstream DCs². The typical scenario for this activity is a DC gathering data for the convenience of an AC at the same institution. We allow for this practice, but the data obtained by a pull mechanism becomes outside the IGS data flow.

DCs must regard the pushed data as READ-ONLY. The only exception to this rule is that a DC may agree to accept data in alternate compression formats (e.g., .gz) and recompress the files in the IGS standard compression mode (.Z, Unix compression) prior to publishing. The underlying RINEX file must still be regarded as READ-ONLY.

All data transfers must be verified as uncorrupted, per IGS required guideline #2.1.17.³ Failed transfers should be retried at intervals until successful.

5. Examples

Examples are given below. In the text following, “from AAAA” is understood to mean “from AAAA’s OC.”

Given this setup table:

OC -----	primary DC -----	secondary DC -----
AAAA	RDC1	GDC1
BBBB	RDC1	RDC2
CCCC	GDC1	GDC2
DC -----	primary DC -----	secondary DC -----
RDC1	GDC1	GDC2

² Exceptions to this rule will be considered, but only to satisfy practical constraints on the development resources of participating DCs; controlled testing of the strategies outlined herein shall determine their technical soundness and the practicality of their implementation.

³ See IGS Site guidelines, via <http://igscb.jpl.nasa.gov/network/netindex.html>

we can distinguish three different cases:

- AAAA uploads to an RDC(primary) and GDC (secondary)
- BBBB uploads to two RDCs
- CCCC uploads to two GDCs

Data flow from the RDC to the GDC is defined in the second part of the table. Using the setup table, the data flow of each station is clearly defined.

5.1 Upload to a RDC (primary) and GDC (secondary)

In Figure 6, we have detailed the data flow for AAAA, which is uploading its data to RDC1 (primary) and GDC1 (secondary). The principal data path is AAAA->RDC1->GDC1 and from there to the other GDCs on the gdcmirror channel. If this path is completed, GDCs know to discard copies received on the backup channel from AAAA or RDC1.




Figure 6: Data flow example where station uploads to an RDC and GDC

Should RDC1 (primary DC) be disabled, data flow continues from AAAA to GDC1 (secondary DC) on the backup channel, and from there (following a 5 minute wait) to the other GDCs on the gdcmirror channel. When RDC1 becomes available, it will either A) receive and archive a copy from AAAA (which retried at regular intervals to upload its data) or B) collect a copy of AAAA from a GDC, then transmit that to GDC1 and GDC2. In both instances the GDCs will discard this copy since have already received this data.

If a file replacement is necessary, then AAAA must use the replacement channel to send a replacement file to RDC1 and GDC1. GDC1 will then transmit its copy to the other GDCs on the gdcmirror channel; RDC1 will forward its copy to GDC1 and GDC2 on the replacement channel. It is evident that a replacement results in a small storm of activity as GDC1 and GDC2 also push these copies to the other GDCs. We judge this to be acceptable in light of the fact that equivalence is guaranteed across GDCs, and replacements are to be the exception rather than the norm.

5.2 Upload to two RDCs

For BBBB uploading to two RDCs: RDC1 (primary) and RDC2 (secondary), the scheme in Figure 7 is obtained. The primary data path is BBBB->RDC1->GDC1 and from there to the other GDCs on the gdcmirror channel.

If this path is completed, GDC2 will discard a second copy sent from RDC2, as well as a copy sent on the backup channel from RDC1.

If GDC1 is disabled, GDCs 2,3, and 4 will receive the BBBB data via RDC1->GDC2 or RDC2->GDC2, whichever is completed first. When GDC1 is available for uploads again, it may receive copies of BBBB from either RDC1 or another GDC.




Figure 7: Data flow example with uploads to two RDCs.

5.3 Upload to two GDCs

Finally, Figure 8 displays the comparatively simpler data flow of CCCC, which is uploading to GDCs: GDC1 (primary) and GDC2 (secondary).




Figure 8: Data flow example with uploads to two GDCs.

6. File revisions

The remaining question is: how do the end users know that a file has been replaced? Here are some ideas.

1. When a file is revised, the OC must generate an email message with a standardized Subject (e.g. “AAAA 2006 140 revised due to missing data.” The message is sent out on a mailing list, and ACs can use either humans or scripts to read the messages and act accordingly.
Problem: What if the OC fails to generate the message?
2. The OC must put a COMMENT in the RINEX header with the reason for replacement, and the DCs must extract and log the COMMENT. The log could be emailed out when appended to.
Problem: (1) What if the OC fails to generate the COMMENT?
(2) Some additional load on DCs to uncompress and read files.
Note: The comment certainly should be required on general principle.
3. The OC must transmit a separate small file with a standardized line explaining the replacement. DCs will log the replacements. The log could be emailed out when appended to.
Problem: Again, the OC might fail to generate the file. Then what?
4. Whenever a DC receives a file on the replacement channel or performs an overwrite with a copy received on the gdcmirror channel, it will log the replacement. This might be combined with #2 or #3 to utilize the reason if supplied by the OC, or “Unknown” if not. The log could be emailed to subscribed users and archived at the DCs when modified.
5. The replacement upload “area” at all GDCs is maintained with an “expiration” interval implemented on the filesystem itself. Read only access may be granted, through another account

at the same DC, to this area, providing any user with the option to “scan” this area and immediately recognize which IGS station data files have been replaced recently. The interval would be identical at all GDCs and no file would be removed from this area until the age of the timestamp on the file has exceeded this interval.

A final note: These rules might also be applied to products to ensure that product replacements are handled and archived properly.

7. Thoughts on Implementation

More detailed steps will need to be discussed with the IGS DCs in order to implement the concepts detailed in this position paper. These steps should be implemented in an incremental fashion, with a clearly defined objective and timeline for each step, and shall include (minimally) the following:

- Gather information on present station/OC data flow
- Define the two upstream DCs for each station/OC and update IGS site logs with this information ; the DCWG chair will examine the proposed primary & secondary DCs for sensibility and to avoid "routing loops"
- Prepare procedures at all OCs/ODCs for upload to two DCs at higher level
- Prepare procedures at all RDCs and GDCs to receive data from two upload paths and distinguish principal, backup, replacement and (for GDCs only) gdcmirror channels
- Test data flow procedures thus verifying operational, backup, and replacement data flow
- Document data flow and provide information at IGS CB

The authors assume the objectives set forth in this plan may not be fully recognized or achieved before a significant amount of time – possibly one or more years – has passed, in order to minimize the burden of implementation on the IGS community as a whole.

8. Recommendations

1. Flow of IGS station data from station/OC to a higher-level data center is performed using a PUSH ONLY set of rules, with the ultimate objective being “All IGS station data available, as copies (referring to content, not necessarily compression, filesize and final checksum), to the public at all IGS GDCs as soon as possible”.
2. Stations/OCs, ODCs, and RDCs will define primary and secondary data centers to push their data to and will update IGS site logs with this information
3. Stations/OCs should document replacement of data files and notify the IGS through automated procedures

Acknowledgement

A. Moore's portion of this research was performed at the Jet Propulsion Laboratory, California Institute of Technology, under a contract with the National Aeronautics and Space Administration.