

Response to reviewer #1

1. General comments:

- On some parts of the paper and while talking about storage, the Big-O notation is used (but this is not consistently used through the work). I believe that given the tone of the paper, this notation is not needed and just mentioning the units (e.g. 20TB) is enough.
 - The Big-O notation has been removed throughout.
- Data integrity validation (after data distribution) is not explicitly mentioned on the paper (even though it can be inferred from some of the statements). Could you please add some specific text on how you ensure that the distributed and/or replicated data has not been corrupted or tampered in any way?
 - Data integrity is ensured through file checksumming and is a core component of the data transfer tool Synda, the checking has been described in the section 3.3.5 on Synda.
- Disaster recovery is not mentioned in this work: I think it is an interesting aspect to discuss, especially given that the data is sharded. Could you please add some information about this?
 - It is the responsibility of each modelling centre to ensure that their data is backed up (usually via a tape archive), this is not the primary goal of the ESGF. However, where replica data exists there is a certain amount of redundancy in the system. It is usually the case that the large data centres that obtain replica copies have automated archive backup copies as part of their normal workflow, additionally the replica copies themselves can act as a form of backup. The publishing data node should retain a copy of the metadata as best practice for any production level database, however were there to be a problem it theoretically possible for metadata to be restored from the shards, in practice this has never been necessary as during any software or hardware updates/migrations sites have always ensured best practice and maintained their own backup copies of the metadata catalogues that underpin the ESGF search. Statements on this have been incorporated on p5, para. 2.

2 Specific Comments:

- P2, 143: Could you please reference the recommended security policies?
 - The sentence to which this refers was slightly ambiguous, it was referring to the security polices of the ESGF software stack. The documentation for this (<https://esgf.llnl.gov/esgf-media/pdf/ESGF-Software-Security-Plan-V1.0.pdf>) has been included as a footnote along with a change in the bullet point for clarity. Update of system software has been noted in bullet point 2, however the main focus here is on the ESGF and value-added service software.
- P3, 145: Could you please reference the necessary policies?

- Policy document <https://esgf.llnl.gov/esgf-media/pdf/ESGF-Policies-and-Guidelines-V1.0.pdf> has been included as a footnote.
- P3, 146: Could you please reference the required resources?
 - The resource requirement documentation <https://esgf.llnl.gov/esgf-media/pdf/ESGF-Tier1and2-NodeSiteRequirements-V5.pdf> has been included as a footnote.
- P4, 177: "federated data archive": Could you please add a reference to the architecture (if available)?
 - This is covered by the references to Williams et al 2011,2016 and Cinquini et al 2014 and have been reiterated.
- P4, 188: "standardised set of rules": Could you please reference these rules?
 - The sentence referred to here is misleading and has been changed to read "standardised terminology", a reference to the CMIP6 controlled vocabularies (https://github.com/WCRP-CMIP/CMIP6_CVs) containing the standardised terms is included as a footnote.
- P4, 190: "HTTP": Is TLS supported? If so, could you please state it?
 - Yes TLS is supported as best security practices for SSL config are implemented, this has been stated.
- P5, 191: "in an automated way": Could you please add more details and/or references about the automations?
 - The automation is managed by Solr, this had been described.
- P5, 192: "see the same search results from any index node". How is split-brain and replication latency resolved here? (Maybe you can add some more context on the paper).
 - The federation works toward eventual consistency. Updates at one site (new publications) won't be available instantly at all sites. On occasion there might be delays due to synchronisation of the Solr shards but these will only last a couple of minutes at most, this latency is not a problem in this context. There is no shared state that needs to be maintained among index sites that requires sophisticated consistency protocols.
- P5, 1103: "18PB": Previously you stated 20PB, is this a more accurate result or maybe a typo?
 - There are a few different estimates on the final size of CMIP6, the estimate 20PB is the most general and should be used at this stage, this has been corrected.
- P7, 1140: "includes a database": What database (technology) is this? Maybe you can also put a reference.
 - This is an XML database and is described in the reference as given in the manuscript.
- P10, 1110: I think it could be interesting some details on how the PID (the id value itself) is formed (also maybe some reference on how it is calculated to be unique and avoid clashes).
 - The PID is of the form hdl:21.14100/<uuid> where uuid is the unique id generated by the CMOR uuid library (using the general open source uuid library), see https://www.earthsystemcog.org/site_media/projects/wip/CMIP6_PID_Implementation_Plan.pdf

- P12, 1284: Are you using Elastic Search as part of the analyzer? If so, maybe you could state it. Otherwise, could you please add some references or links to the code of the analyzer.
 - No, we are not using Elastic Search. The analyzer acts as an ETL (Extraction-Transformation-Loading) and is composed of a set of scripts written in Python and interacting with an OLAP database. The Elastic Search components are only used for logs shipping & collection. The analyzer code is a part of the Dashboard code a link has been included as a footnote.
- P12, 1284: About the dashboard: For what I can see you are using a custom-built UI. Any reason not to use Kibana (as you are already using part of the ELK stack)? If so, I think it is relevant to get it mentioned here. Also I think it is interesting to link the code repository for the dashboard in here.
 - We are not using Kibana as Kibana is designed to use Elasticsearch as a data source which we are not using. Our UI does not query any Elastic Search database, but rather a set of data marts (OLAP-driven approach) created through an ETL procedure consisting of a set of Python scripts. A link to the dashboard repository on GitHub has been included.
- P13, 1288: "recommended a basic hardware": Could you please list these hardware requirements or put a reference to them?
 - The recommended hardware configurations are described in the reference Dart 2013, this has been reiterated.
- P14, 1309: After reading all the section 5 it seems you are describing a (somehow manual) deployment procedure. Did you follow any specific methodology, if so, could you please mention it and maybe add a deployment diagram (it will add more clarity to all the section 5).
 - This procedure was indeed mostly manual no specific methodology was applied. The aim was simply to perform a series of whole system tests that increased in complexity at each iteration. While the installation of some software such as the ESGF software stack step (1) would be considered deployment, the other steps are not related to deployment rather they are operational usage. The procedure listed was performed in a strictly linear fashion and so I do not see how a deployment diagram would be beneficial.
- P17, 1383: "to run the test suite": Can you please put a reference (if available) to the test suite code and/or documentation?
 - A link to the test suite code repository has been included.
- P17, 1435: "In order to assist sites that were not able to participate to effectively publish": Are there any plans to automate all the steps described in the document? If so, could you please state it on the paper?
 - There are no plans to automate this procedure as too many of the steps are manual, particularly requiring intervention when problems/conflicts arise. However, you raise an important point, as the underlying software and architecture continues to evolve methods for more automated testing should be considered a priority. Unfortunately, due to the nature of this it is likely impossible to automate all of the steps but greater emphasis should be placed on this in the future. It should be noted that each site operates somewhat independently and the larger sites that have more resource and experience do have some automated workflows, we should consider how best to share these approaches bearing in mind that there is no one size fits all solution.

Anonymous Referee #1

1. General comments
 - a. Big-O notation has been removed throughout.
 - b. Data integrity is ensured through file check summing and is a core component of the data transfer tool Synda, the checking has been described in the section on Synda lines 268-271
 - c. Disaster recovery
2. Specific comments
 - P2, 143: Could you please reference the recommended security policies?
 - P3, 145: Could you please reference the necessary policies?
 - P3, 146: Could you please reference the required resources?

 - P4, 177: "federated data archive": Could you please add a reference to the architecture (if available)?
 - DONE
 - P4, 188: "standardised set of rules": Could you please reference these rules?

 - P4, 190: "HTTP": Is TLS supported? If so, could you please state it?

 - P5, 191: "in an automated way": Could you please add more details and/or references about the automations?

 - P5, 192: "see the same search results from any index node". How is split-brain and replication latency resolved here? (Maybe you can add some more context on the paper).
 - P5, 1103: "18PB": Previously you stated 20PB, is this a more accurate result or maybe a typo?
 - DONE

 - P7, 1140: "includes a database": What database (technology) is this? Maybe you can also put a reference.
 - it is an XML database. The reference is as given in the text you've quoted ..I'm not sure what else is needed there.

 - P10, 1110: I think it could be interesting some details on how the PID (the id value itself) is formed (also maybe some reference on how it is calculated to be unique and avoid clashes).
 - Martina/Stephan

the handle pid is structured according to:

hdl:21.14100/<uuid>

14100 is the globally registered domain for CMIP6 PIDs and <uuid> is the unique id generated by the CMOR uuid library (using the general open source uuid library)

see

also https://www.earthsystemcog.org/site_media/projects/wip/CMIP6_PID_Implementation_Plan.pdf:

“21.14100” is a Handle prefix that is administrated by DKRZ as part of a contract with GWDG, a newly established Multi-Primary Administrator within the framework of DONA. DONA (Digital Object Numbering Authority) is a Swiss foundation which has taken over administrative control of the Handle System from CNRI. DKRZ manages the future prefix namespace of 21.14xxx”

- P12, 1284: Are you using Elastic Search as part of the analyzer? If so, maybe you could state it. Otherwise, could you please add some references or links to the code of the analyzer.
 - No, we are not. The analyzer acts as an ETL (Extraction-Transformation-Loading) and is composed of a set of scripts written in Python and interacting with an OLAP database. The Elastic Search components are only used for logs shipping & collection.
- P12, 1284: About the dashboard: For what I can see you are using a custom built UI. Any reason not to use Kibana (as you are already using part of the ELKstack)? If so, I think it is relevant to get it mentioned here. Also I think it is interesting to link the code repository for the dashboard in here.
 - We are not using Kibana (“Kibana is designed to use Elasticsearch as a data source”, from the Kibana website). Our UI does not query any Elastic Search database, but rather a set of data marts (OLAP-driven approach) created through an ETL procedure consisting of a set of Python scripts.
- P13, 1288: “recommended a basic hardware”: Could you please list these hardware requirements or put a reference to them?
- P14, 1309: After reading all the section 5 it seems you are describing a (somehow manual) deployment procedure. Did you follow any specific methodology, if so, could you please mention it and maybe add a deployment diagram (it will add more clarity to all the section 5).
 - This procedure did require a lot of manual deployment. I am not sure how useful a deployment diagram would be given the serial nature of the procedure.

- P17, 1383: "to run the test suite": Can you please put a reference (if available) to the test suite code and/or documentation?
 - Add the link <https://github.com/ESGF/esgf-test-suite>
- P17, 1435: "In order to assist sites that were not able to participate to effectively publish": Are there any plans to automate all the steps described in the document? If so, could you please state it on the paper?
 - There are no plans to automate this procedure as too many of the steps require manual intervention, particularly if there is a problem. This is an interesting question however as the underlying architecture changes in the future and becomes ever more complex. Should the same infrastructure be used in a further CMIP round this process should be designed with as much automated deployment and testing as possible.