

Manuscript Number:	GIGA-D-20-00202R1	
Full Title:	Transcriptome annotation in the cloud: complexity, best practices and cost.	
Article Type:	Technical Note	
Funding Information:	U.S. National Library of Medicine (Intramural Research Program of the National Library of Medicine, National Center for Biotechnology Information at the National Institutes of Health.)	Dr. David Landsman
Abstract:	<p>Background</p> <p>The NIH Science and Technology Research Infrastructure for Discovery, Experimentation, and Sustainability (STRIDES) initiative provides NIH-funded researchers cost-effective access to industry-leading commercial cloud providers, such as Amazon Web Services (AWS; Seattle, WA, USA) and Google Cloud Platform (GCP; Mountain View, CA, USA). These cloud providers represent an alternative for the execution of large computational biology experiments like transcriptome annotation which is a complex analytical process that requires the integration of multiple biological databases and several advanced computational tools. The core components of annotation pipelines published since 2012 are BLAST sequence alignments using annotated databases of both nucleotide or protein sequences almost exclusively with networked on premises compute systems.</p> <p>Findings</p> <p>We present a comparative study of multiple BLAST sequence alignments using two public cloud providers: AWS and GCP. We have prepared several Jupyter Notebooks with all the code required to submit computing jobs to the batch system on each cloud provider. We consider the consequence of the number of query transcripts in input files and the effect on cost and processing time. We tested compute instances with 16, 32 and 64 vCPUs on each cloud provider. Four classes of timing results were collected: the total run time, the time for transferring the BLAST databases to the instance local solid state disk drive (SSD), the time to execute the Common Workflow Language (CWL) script and the time for the creation, setup and release of an instance. This study aims to establish an estimate of the cost and compute time needed for the execution of multiple BLAST runs in a cloud environment.</p> <p>Conclusions</p> <p>We demonstrate that the public cloud providers are a practical alternative for the execution of advanced computational biology experiments at low cost. Using our cloud recipes, the BLAST alignments required to annotate a transcriptome with ~500,000 transcripts can be processed in less than 2 hours with a compute cost of about 200-250 USD. In our opinion, for BLAST based workflows, the choice of cloud platform is not dependent on the workflow but, rather, on the specific details and requirements of the cloud provider (e.g. NCBI maintains updated copies of the very large genetic sequence databases, such as nr, RefSeq and SRA, on both GCP and AWS). These choices include the accessibility for institutional use, the technical knowledge required for effective use of the platform services, and the availability of open-source frameworks such as application programming interfaces (APIs) to deploy the workflow.</p>	
Corresponding Author:	David Landsman, Ph.D National Center for Biotechnology Information Bethesda, Maryland UNITED STATES	
Corresponding Author Secondary Information:		
Corresponding Author's Institution:	National Center for Biotechnology Information	
Corresponding Author's Secondary Institution:		
First Author:	Roberto Vera Alvarez, Ph.D	

First Author Secondary Information:	
Order of Authors:	Roberto Vera Alvarez, Ph.D
	Leonardo Mariño-Ramírez, Ph.D
	David Landsman, Ph.D
Order of Authors Secondary Information:	
Response to Reviewers:	<p>Reviewer reports:</p> <p>Response: We thank the referee for the comments and for their time dedicated to this manuscript.</p> <p>Reviewer #1: In this manuscript, the authors describe the usability and the performance of public cloud computing services and their batch job execution services for the transcriptome annotation pipeline. The authors also compare the services of two major public cloud vendors, Google Cloud Platform (GCP), and Amazon Web Service (AWS). The results show that the two cloud providers can provide a similar experience on its operation, workflow execution time, and payment for execution. The performed experiments follow the modern best practice of data analysis using tools and frameworks for better reproducibility, including Docker container, the Common Workflow Language (CWL), and Jupyter Notebook. This article can be a good example of a reproducible study with open data and open-source software.</p> <p>This report is very significant with the practical statistics, which can be a helpful reference for all the cloud use cases in biomedical data analysis. The title states that this study focuses on the transcriptome annotation, but the output provides insight for all the cloud use cases. I suggest the authors change the title because the current one looks the best practice valid only for the transcriptome annotation in the wide variety of genomic data analysis.</p> <p>Response: We thank the referee for their comment and suggestions. It is true that the results of this study provide insight for many computational biology workflows but our experiments were limited to the transcriptome annotation process, specifically to the BLAST searches that are the core of the annotation. We think that the best practices and conclusions of this study should remain limited to the cloud transcriptome annotation process.</p> <p>Below are minor comments for the manuscript.</p> <p>Page 3, the last paragraph: The authors claim as "little has been published describing cloud costs and implementation best practices". However, there is a study implemented software to monitor the runtime metrics of a given workflow, and support the cost estimation for the executions on the cloud (https://doi.org/10.1093/gigascience/giz052). This article describes the cost for the normal EC2 instance and does not mention the batch execution services, yet it provides additional information to the readers. Please consider introducing this in the background section as a related study. (Disclaimer: I am the first author of this article)</p> <p>Response: We agree with the referee that this study provides additional information about the use and cost of computational biology workflows in the cloud. A new paragraph was added to the Background section: "The utilization of cloud environments for computational biology experiments is increasing [14-17], however, little has been published estimating cloud costs and implementation best practices. A recent work published by Ohta at al. [18] presents a tool named CWL-metrics that collects runtime metrics of Docker containers and workflow metadata to analyze workflow resource requirements. This study presents a cost estimation for the execution on the cloud for AWS EC2 instances, but does not mention the cloud batch system for users to submit thousands of jobs to the cloud."</p> <p>Page 12, Best practices The authors describe here that they recommend CWL as a workflow language. I think the authors should introduce the reason they chose CWL at the Method section where</p>

they first mention CWL (Page 4, Transcriptome Annotation Workflow). The authors also should provide a more practical reason to choose CWL because CWL is not only one framework that has portability and scalability. For example, having multiple workflow runners or its syntax easily parsed and connected to the resources such as runtime metrics can be a reason to choose CWL in this use case.

Response: We agree with the referee about the inclusion of more information about CWL and the reason why it was selected for our study. We added a new section to the manuscript named: Common Workflow Language

Workflow dependencies

Both AWS batch and Google LifeScience allow users to specify only one container per one batch job. This means the user needs to install the workflow runner (in this article cwltool) and the tools to process the data (e.g., BLAST) in a single container. However, the current best practice for bioinformatics is to separate the containers for each tool. Thus, in most cases, users cannot reuse the containers they use for a normal computing environment to the cloud batch system. The authors need to mention the limitation that one needs to create a container for cloud batch systems. Another possible solution would be to run sibling containers from the main batch job container, though I do not know if it is feasible with AWS and GCP.

Response: The referee's comment is correct. Docker-in-Docker, the process to execute docker containers inside of another Docker container is not allowed in both GCP and AWS. We add a new section to the manuscript named GCP and AWS batch system limitations to address these limitations.

Software and framework related to the cloud batch services Many workflow languages and runners are supporting AWS batch and GCP. For example, Nextflow (<https://www.nextflow.io/docs/latest/awscloud.html>), Cromwell (<https://cromwell.readthedocs.io/en/stable/backends/Google/>), or Snakemake (<https://snakemake.readthedocs.io/en/stable/executing/cloud.html#executing-a-snakemake-workflow-via-tibanna-on-amazon-web-services>) can run the workflows via AWS batch or GCP. Task Execution Service (TES) of the Global Alliance for Genomics and Health (GA4GH) Cloud working group is also a framework to utilize the cloud batch system (<https://github.com/ga4gh/task-execution-schemas>). Some tools that run workflows on the cloud using the ETL framework (<https://doi.org/10.21105/joss.01069>) or cloud batch services (<https://github.com/DataBiosphere/dsub>) are also available. These may not be directly relevant to this study, but it would be helpful for readers to understand how the cloud batch services are being used by the researchers.

Response: We agree with the referee about the inclusion of more information about workflow managers. We added a new section to the manuscript named: Common Workflow Language.

Reviewer #2: Reviewer's report

Title: Transcriptome annotation in the cloud: complexity, best practices and cost. transcriptome data

Reviewers: Qiao Xuanyuan and Lucas B. Carey

Response: We thank the referees for their comments and time dedicated to this manuscript.

Reviewer's report:

The authors provide a perspective on transcriptome annotation in the cloud by presenting a comparative study of the two main public cloud providers, aiming to help the reader determine the proper cloud platform and its utilization in their research. The main strength of this paper is that it addresses a question that little has been studied before—the cloud cost estimates and implementation best practices. This is useful not only for labs doing transcriptome annotation, but, because all code is provided and very well commented, it might be of use for labs dealing with big data & distributed computation problems in general.

The manuscript is well written, and I have only a few minor concerns on presentation and the information that is provided.

1. The tested transcriptome data need to be clarified. I read the Data partitioning code for the creation of 20 FASTA files of the different query sizes (Fig 3). The variation in processing time among these files is consistent across clusters, and is therefore presumably due to differences in query sets in each file. This is surprising, as 10,000 is a large number. Are the differences because the sequence records were created sequentially from the input fasta file, with some genes having large numbers of hits? Would subsetting transcripts randomly result in more uniform processing times? There is almost two-fold variation in processing time between query files.

Response: The referee's comment is correct. The variation in processing time for queries with the same number of transcripts was due to the sequentially subsampling approach. We have modified the Data Partitioning notebook to create a random sampling of the transcriptome file. The notebook now shows the statistics for each file created. All measured parameters, mean, standard deviation, minimum length, 25%, 50% and 75% quarters show similar values for all files. Modified text was added in the beginning of the Results and Discussion section "From the Opuntia pool of transcripts, we analyzed three sizes of query files: 2,000, 6,000, and 10,000 transcripts in each input query file. Two experiments were executed. First 20 FASTA files (input files for the workflow) for each query size were randomly created, see notebook "01 - Data Partitioning". Each of these files were submitted independently as jobs to the batch systems on each cloud provider. For the second experiment, 120,000 transcripts were randomly selected and then partitioned in files with 2,000, 6,000, and 10,000 transcripts to analyze the relationship between query size, runtime and cost."

2. Please include a figure showing the distribution of transcript lengths for *Opuntia streptacantha*, and write that it is the prickly pear cactus. Presumably timing depends on the transcript lengths and on the number of BLAST hits.

Response: The referee's comment about that the timing certainly depends on the transcript length and the number of BLAST hits. We added a cell to the 01 - Data Partitioning notebook that shows the transcriptome length distribution and its statistic metrics.

3. It is difficult to draw the conclusions from the way the data are plotted in Figure 4. The author concluded that "Reducing the number of transcripts per input file will reduce the total running time but will also increase the cost of the project as more instances will be in used."

In addition to the raw data boxplot, it is better to show how the time and cost scale with query size, or with the number of instances. (The number of instances equals query size divided by total transcripts). Controlling the total transcripts and making instances as variable may be helpful in balancing the interpretation and data.

The plot below shows the relationship between the number of CPUs and time, and query size and time. (data collected using <https://automeris.io/WebPlotDigitizer/>). It also provides direct-viewing evidence to support the author's conclusion "AWS is more suitable for large data analysis groups to establish a set of queues and compute environments for multiple pipelines." Unlike the boxplot in the current manuscript, this figure also shows the differences in scaling between 16, 32 & 64 vCPU nodes.

Please add graphs showing query file sizes vs time (as below) and query file sizes vs cost. As well as cost vs time. These are the important take-home messages from the manuscript, but it is difficult to extract this information from the current figures.

Response: We agree with the referee's comment that the relationship between query size, time and cost was not well described in the way that the results were presented in Figure 4. We add a new experiment where we processed 120,000 transcripts using the three query sizes. The new Figure 4 shows the relationship between the time, cost and query size for processing a fixed number of transcripts. Additional text describing this experiment was added to the Result and Discussion section: "Figure 4 shows the time and cost of processing 120,000 transcripts using second generation 64 vCPUs instances on each cloud provider. Reducing the number of transcripts per input file reduces the total run time but will also increase the cost of the analysis as more instances will be used. BLAST databases are transferred to more instances spending, on average, 10 minutes for each instance. For example, our experiment with the

10,000-query size processes all transcripts in about 105 minutes with a total cost of 59.37USD using 12 instances (GCP, N2, 64 CPU). Processing the same number of transcripts with a query size of 2,000 costs 122.36USD with all transcripts processed in 43 minutes using 60 instances (GCP, N2, 64 CPU). “

Optional (not necessary) suggestions:

The figures could be improved to give a clearer visualization of the data. For Figure 3a, will adding a secondary Y-axis regarding money and draw a line plot be better to show the relationship between total time and cost? For Figure 3b&c&d, drawing a component bar chart could allow the readers to compare the job-dependent time among various configurations and demonstrates the proportionality at the same time.

Response: We thank the referee for this suggestion. In our opinion, as the data we are plotting here is not continuous data, drawing a line between points in Figure 3a could imply the idea of continuity. For Figures 3b, 3c, and 3d drawing bars could hide the variability of the runtime for each input file that is associated with instance and network performance. We prefer to keep the figure as it is now.

Reviewer #3: The authors provide a comparison of two cloud-based solutions for running BLAST-based transcriptomics analysis.

With cloud-based solutions becoming more popular in science, I think this comparison, along with the practical recommendations provided in this manuscript will be interesting to readers.

Response: We thank the referee for their comments and time dedicated to this manuscript.

Some suggestions for enhancements below:

1) The authors mention that there are numerous genomics companies in the space of cloud based biocomputations, citing reference 14 which discusses some of the legal responsibilities of groups doing such cloud based analyses. However, the authors do not connect this discussion back to the use of GCP or AWS, where users would need to obtain similar guarantees of data security which may not be possible to obtain. Thus a comparison of pricing against the private firms which provide similar services would be interesting to see if they provide specific guarantees that affect researchers with data that requires specific legal requirements.

Response: We thanks the referee for the comment. This study was executed under the NIH's STRIDES initiative using the currently available cloud provider partners: AWS and GCP. We do not consider that a comparison between public and private cloud providers and their legal responsibilities is within the scope of this study. However, we think that mentioning some of the most important private cloud providers is necessary. Accordingly, we rephrased the paragraph as:

In addition, private genomic cloud providers, for instance DNAnexus (www.dnanexus.com), DNASTAR (www.dnastar.com), Seven Bridges (www.sevenbridges.com) and SciDAP (scidap.com), also are in the market and offer cloud-based genomics frameworks. These commercial cloud providers make the execution of computational biology experiments easier offering command line and web-based interfaces designed for genomic data analysis.

2) I really like the inclusion of a best practices section with practical recommendations, and would love to see this section expanded:

a) In the first point the authors state: "We recommend CWL because the resulting product is portable and scalable, and it can be executed across a variety of computational environments as dissimilar as personal laptops or the cloud". However, these features are not exclusive to CWL, and solutions such as NextFlow and SnakeMake (and probably others) would also fit this description (and both of these also offer point 2 Conda and containerization). Please elaborate on your recommendation to include discussion of other workflow management systems, and explain in more

detail why you would recommend CWL over these other solutions.

Response: We agree with the referee about the inclusion of more information about CWL and workflow managers. We added a new section to the manuscript named: Common Workflow Language.

b) In point 5, you recommend that users "Execute a small test in the cloud to find the best instance type for a workflow". Do you have any further practical recommendations about how users can best go about this? E.g. how does one define a "small test run" from a full datasets, and how can they predict how this will scale up to the full analysis and determine the most suitable machine types?

Response: We agree with the referee that defining a "small test run" may be difficult and it is intrinsically determined by the data and the workflow to be used. Our intention with this recommendation was to alert users that the cloud is a completely different environment than local workstations or on premise clusters. Users should test different cloud services and configurations before submitting a huge number of jobs. We edited that recommendation to:

5. Cloud computing behaves differently than local workstations or on premise clusters. Users should define and execute small tests with their data and workflow before submitting large jobs. Testing different cloud services and configurations could help to reduce the runtime and cost for the whole analysis.

3) It could be nice to expand the section about the Jupyter notebooks, and how these are being used, perhaps with some screenshots of results, and/or a small schematic showing that (if I understand correctly): the user interacts with the Jupyter notebook on their local machine, which in turn configures the cloud resources and starts the CWL workflow on the cloud, and then fetches the relevant results back and analyses them and displays results to the user. I think a schematic to this effect would be helpful for less technical readers and this in combination with some screenshots of the analysis results in Jupyter will increase the appeal of your work to research scientists.

Response: We agree with the referee about expanding the Jupyter notebook section. We added more description to it. The notebooks are available on Github for reading and browsing. Adding more figures to the manuscript would increase its size and complexity. Jupyter notebooks are very popular and we think that less technical readers could easily find documentation about Jupyter notebooks without any problem.

4) In the conclusion the authors state "In our opinion, the choice of a cloud platform is not dependent on the workflow but, rather, on the specific details of the cloud provider". However, I don't believe the authors can make this statement having tested only a single workflow. So please rephrase the conclusion, or compare performance of different workflows covering a range of different characteristics (e.g. one that is memory-intensive, one that is CPU-intensive, and one that requires a lot of data transfer) and showing whether this conclusion holds, or whether some of the "specific details of the cloud provider" may make it more or less suitable for certain types of workflows.

Response: We agree with the referee that the phrase was general. We rephrased to this:

In our opinion, for BLAST based workflows, the choice of cloud platform is not dependent on the workflow but, rather, on the specific details and requirements of the cloud provider (e.g. NCBI maintains updated copies of the very large genetic sequence databases, such as nr, RefSeq and SRA, on both GCP and AWS). These choices include the accessibility for institutional use, the technical knowledge required for effective use of the platform services, and the availability of open-source frameworks such as application programming interfaces (APIs) to deploy the workflow.

5) AWS, Google, and Azure are probably the "big 3" providers that most readers will know about, and they might wonder why Azure was not included here and how it would compare. I understand that the authors cannot compare all providers, but it may be useful to at least mention Azure in the introduction where different providers are

	<p>mentioned, and briefly explain if there were any specific reasons why you chose to compare AWS and GCP, and whether the same methodology could also be applied to Azure.</p> <p>Response: We thank the referee for the comment. This study was executed under the NIH's STRIDES initiative using the current available cloud provider partners: AWS and GCP. We don't have access to Azure, therefore, we cannot extrapolate the conclusions of this study to Azure or any other cloud provider.</p>
Additional Information:	
Question	Response
Are you submitting this manuscript to a special series or article collection?	No
<p>Experimental design and statistics</p> <p>Full details of the experimental design and statistical methods used should be given in the Methods section, as detailed in our Minimum Standards Reporting Checklist. Information essential to interpreting the data presented should be made available in the figure legends.</p> <p>Have you included all the information requested in your manuscript?</p>	Yes
<p>Resources</p> <p>A description of all resources used, including antibodies, cell lines, animals and software tools, with enough information to allow them to be uniquely identified, should be included in the Methods section. Authors are strongly encouraged to cite Research Resource Identifiers (RRIDs) for antibodies, model organisms and tools, where possible.</p> <p>Have you included the information requested as detailed in our Minimum Standards Reporting Checklist?</p>	Yes
<p>Availability of data and materials</p> <p>All datasets and code on which the conclusions of the paper rely must be either included in your submission or</p>	Yes

deposited in [publicly available repositories](#) (where available and ethically appropriate), referencing such data using a unique identifier in the references and in the “Availability of Data and Materials” section of your manuscript.

Have you have met the above requirement as detailed in our [Minimum Standards Reporting Checklist?](#)

Transcriptome annotation in the cloud: complexity, best practices and cost.

Roberto Vera Alvarez¹, Leonardo Mariño-Ramírez^{1,2}, and David Landsman^{1,*}

¹ Computational Biology Branch, National Center for Biotechnology Information, National Library of Medicine, NIH, Bethesda, MD, USA.

² Current address: Division of Intramural Research, National Institute on Minority Health and Health Disparities, NIH, Bethesda, MD, USA.

*To whom correspondence should be addressed.

Abstract

Background

The **NIH Science and Technology Research Infrastructure for Discovery, Experimentation, and Sustainability (STRIDES) initiative** provides NIH-funded researchers cost-effective access to industry-leading commercial cloud providers, such as Amazon Web Services (AWS; Seattle, WA, USA) and Google Cloud Platform (GCP; Mountain View, CA, USA). These cloud providers represent an alternative for the execution of large computational biology experiments like transcriptome annotation which is a complex analytical process that requires the integration of multiple biological databases and several advanced computational tools. The core components of annotation pipelines published since 2012 are BLAST sequence alignments using annotated databases of both nucleotide or protein sequences almost exclusively with networked on premises compute systems.

Findings

We present a comparative study of multiple BLAST sequence alignments using two public cloud providers: AWS and GCP. We have prepared several Jupyter Notebooks with all the code required to submit computing jobs to the batch system on each cloud provider. We consider the consequence of the number of query transcripts in input files and the effect on cost and processing time. We tested compute instances with 16, 32 and 64 vCPUs on each cloud provider. Four classes of timing results were collected: the total run time, the time for transferring the BLAST databases to the instance local solid state disk drive (SSD), the time to execute the Common Workflow Language (CWL) script and the time for the creation, setup and release of an instance. This study aims to establish an estimate of the cost and compute time needed for the execution of multiple BLAST runs in a cloud environment.

Conclusions

We demonstrate that the public cloud providers are a practical alternative for the execution of advanced computational biology experiments at low cost. Using our cloud recipes, the BLAST alignments required to annotate a transcriptome with ~500,000 transcripts can be processed in less than 2 hours with a compute cost of about 200-250 USD. In our opinion, for BLAST based

workflows, the choice of cloud platform is not dependent on the workflow but, rather, on the specific details and requirements of the cloud provider (e.g. NCBI maintains updated copies of the very large genetic sequence databases, such as nr, RefSeq and SRA, on both GCP and AWS). These choices include the accessibility for institutional use, the technical knowledge required for effective use of the platform services, and the availability of open-source frameworks such as application programming interfaces (APIs) to deploy the workflow.

Background

The **NIH Science and Technology Research Infrastructure for Discovery, Experimentation, and Sustainability (STRIDES) initiative** (<https://cloud.cit.nih.gov/>) permits NIH supported scientists to explore the use of cloud environments and provides cost-effective access to industry-leading commercial cloud providers. The NIH's STRIDES cloud provider partners, at the time of this study, were Amazon Web Services (AWS; Seattle, WA, USA) and Google Cloud Platform (GCP; Mountain View, CA, USA). Cloud computing offers an *on-demand* model where a user can dynamically allocate “unlimited” compute resources and then release them as soon as the analysis is complete [1]. They offer a reduced cost of compute resources and a friendly user interface that makes cloud computing accessible for large computational biology experiments.

As part of the STRIDES initiative, NIH-funded institutions began to upload and compute data in the cloud. Public biological databases like the Sequence Read Archive (SRA, <https://www.ncbi.nlm.nih.gov/sra/docs/sra-cloud/>) and computational tools like BLAST (https://github.com/ncbi/blast_plus_docs), from the National Center for Biotechnology Information (NCBI), were migrated and are available for public use on AWS and GCP. In addition, NIH-funded researchers are contributing to the NIH's STRIDE initiative not only migrating data analysis workflows to the cloud but also disseminating the suitability of the cloud computing for computational biology experiments.

The annotation of RNA transcripts with functional and biological processes is an important step in developing an understanding of the biological complexity of an organism. Annotation is a challenging process that requires the integration of multiple biological databases and several computational tools to accurately assign a function to an RNA product. Available public information on a target organism is the main limitation of the annotation of non-model organisms. The NCBI Genome database, for instance, contains 54,049 genome-sequencing projects by organism [2]. This includes 12,204 eukaryotes genomes for more than 1,000 species or strains at different assembly levels (95 complete genomes, 1,872 chromosomes, 7,743 scaffolds, and 2,494 contigs (<https://www.ncbi.nlm.nih.gov/genome/browse/#!/eukaryotes/>), accessed on June 30, 2020). Although these data include an important group of organisms, there is a lack of annotation of several species that have significant public health and economic importance. Significantly, in the plant kingdom, *Viridiplantae*, only 3 complete genomes, 331 chromosomes, 625 scaffolds, and 394 contigs are annotated. The advances in next-generation sequencing technologies and the decrease in the cost of sequencing a complete transcriptome is driving a new era in which annotation will be increasing, important and productive.

A review of published manuscripts since 2012 [3-11] reveals that many developed pipelines have a common core component and use the NCBI BLAST tools [12] to align assembled transcriptomes against annotated databases of nucleotides or proteins to identify similarity and

infer function. After an assembly, these alignments are the initial step to identify close and/or distant homologous genes, proteins, and functional domains that could be cross-referenced with other public databases, such as Gene Ontology [13], to generate new annotations of query sequences. As the number of transcripts assembled per study increases, the computing power and storage required to align these transcripts to the BLAST databases also increases. On premises computer infrastructures (including server farms) have been used mainly for the computation of sequence alignments using BLAST. Many laboratories, however, are not equipped with the compute power required for the analysis of increased transcriptome sequencing results. Although a minimum infrastructure could be easy to build and maintain, it may be unnecessary and less financially burdensome with the advent of cloud computing and its utilization in computational biology.

The utilization of cloud environments for computational biology experiments is increasing [14-17]. However, little has been published estimating cloud costs and implementation best practices. A recent work published by Ohta *at al.* [18] presents a tool named CWL-metrics that collects runtime metrics of Docker containers and workflow metadata to analyze workflow resource requirements. This study presents a cost estimation for the execution on the cloud for AWS EC2 instances but does not mention the cloud batch system for users to submit thousands of jobs to the cloud.

Modern cloud providers offer “unlimited” compute resources that can be accessed *on-demand*. An *instance*, as the virtual machines are named in the cloud environment, is deployed using a variety of operating systems like GNU/Linux or Microsoft Windows. Users pay only for the time that the instance is running plus the cost of other resources such as network egress and/or the size of network storage devices. A workflow can be deployed on a manually created instance but this is not cost efficient as the instance will need to be manually reconfigured with workflow dependencies. It will also remain active once the analysis is completed which wastes resources.

Private genomic cloud providers, for instance DNAnexus (www.dnanexus.com), DNASTAR (www.dnastar.com), Seven Bridges (www.sevenbridges.com) and SciDAP (scidap.com), and others, also offer cloud-based genomics frameworks. These commercial cloud providers make the execution of computational biology experiments easier by offering command line and web-based interfaces designed for genomic data analysis.

Most cloud providers offer a batch system that can do the configuration automatically for users to submit several parallel jobs. The batch system makes the process of instance creation, setup and termination fully automatic.

Batch processing is a technique for processing data as a single large collection of iterative steps instead of individually. It reduces user interactivity to process submissions by automating the remaining steps. Modern cloud providers offer a batch system that can be personalized to process many different workflows. Figure 1 shows the component of a generic cloud batch system. It is comprised of a *batch queue* to which users submit the *tasks*. Each *task* uses a *job definition* to create a *job* where all computational resources and the workflow steps are outlined. Then, an *instance* is automatically created with the resources requested by the *job*. Since all the data for the analysis is in the cloud, the *instance* downloads the input data from the *cloud storage system* and, after successfully completing the workflow, uploads the results, releasing all computational resources.

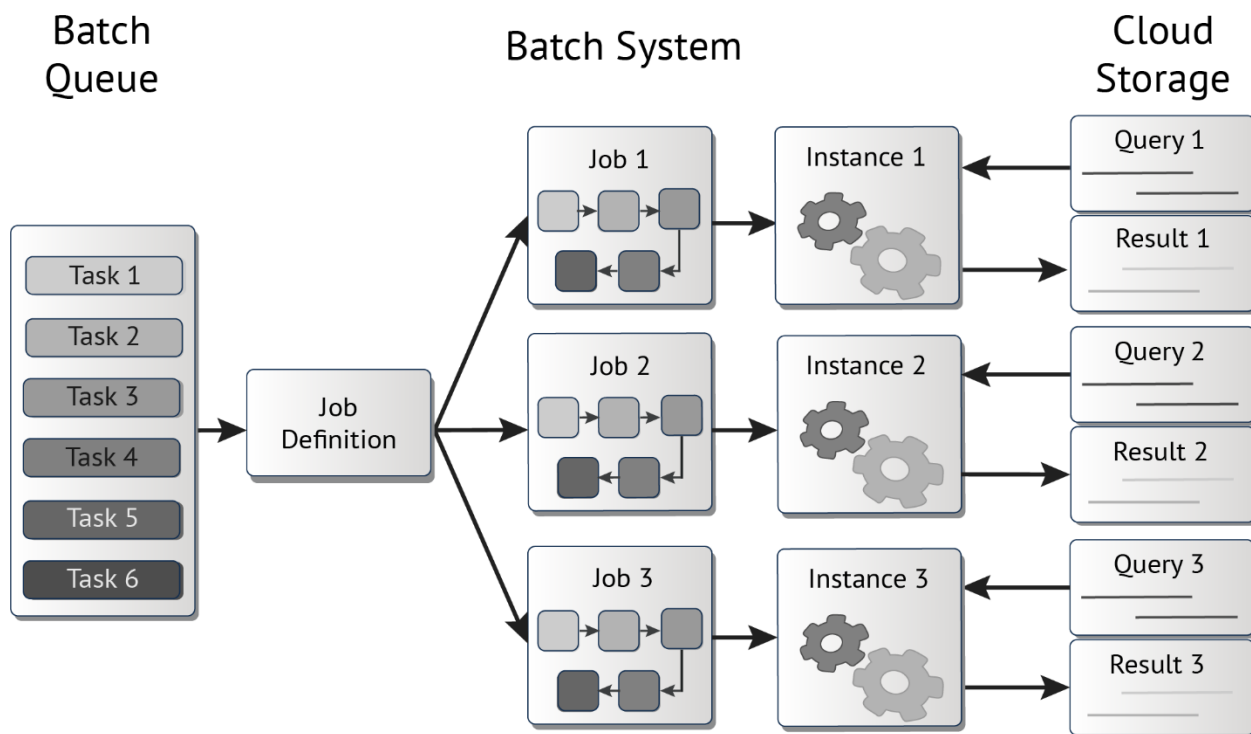


Figure 1: Basic components in a cloud-based batch system

In this manuscript, we present a comparative study of multiple BLAST searches and alignments required to annotate transcriptome data. This study aims to establish an estimation of the cost and time needed for the execution of multiple BLAST searches on the cloud. Our recommendation on best practices for deploying computational biology workflows in the cloud are also presented.

Methods

Transcriptome Annotation Workflow

This study focuses only on the many BLAST alignments which are the most compute-demanding core of a transcriptome annotation process. BLAST alignments require considerable compute resources which generate intermediate results that are used to complete the annotation process. The remaining part of the annotation pipeline is excluded from our study as it can be executed on a workstation and does not require an extensive use of the cloud.

The input for the workflow is a transcriptome in FASTA format. First, TransDecoder [19] is executed to generate all open reading frames (ORFs) from the input file. Then, BLASTP and RPS-BLAST are executed on the TransDecoder output files generating a list of homologous proteins and conserved protein domains (BLASTP uses the BLAST *nr* database, and RPS-BLAST uses the NCBI Conserved Domain Database (CDD) [20]). The transcriptome files are also used as inputs for BLASTN and RPST-BLASTN which are executed using the BLAST *nt* database and the NCBI CDD database, respectively. These processes generate a list of homologous genes and a list of conserved domains, see Figure 2. The workflow was implemented using the Common Workflow Language (CWL) [21] and is freely available at:

https://github.com/ncbi/cloud-transcriptome-annotation/blob/master/bin/cwl-ngs-workflows-cbb/workflows/Annotation/transcriptome_annotation.cwl

The workflow uses as input a FASTA file, which we named *query*, and includes multiple transcripts to be processed. The number of transcripts to be included in a *query* is another parameter that merits an analysis. The size of the *query* affects the workflow processing time as a complete transcriptome could be comprised of thousands to hundreds of thousands of transcripts assembled from a next-generation sequencing (NGS) experiment [22].

Our analysis is based on the execution of the workflow with a batch system provided by each cloud platform. This approach keeps the compute time, and therefore the cost, to a minimum. It also limits the user interaction with the jobs to only the submission step.

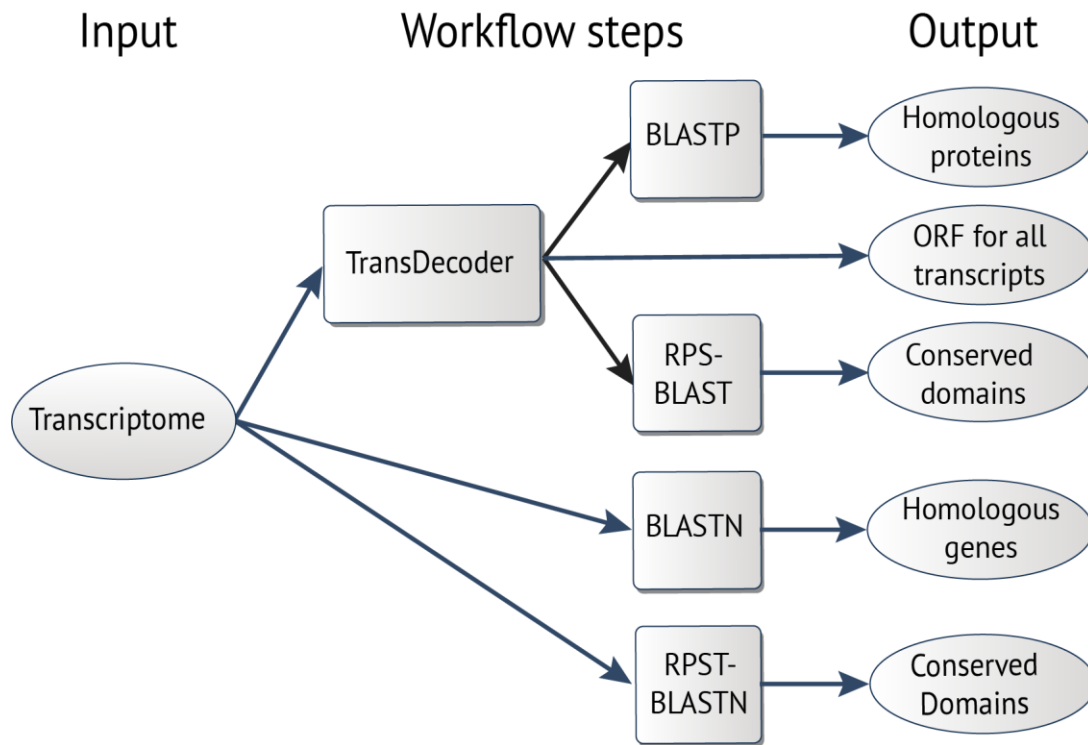


Figure 2: Schema of the transcriptome annotation workflow

Containerized workflows

Containerizing a workflow involves bundling it with all its dependencies and configuration files so that it can be executed across different computing environments. The workflow dependencies in the container uses the same version and compiled libraries when it is executed in any computing infrastructure which would make the process highly reproducible. In this study, we use Docker as the container engine. Docker permits the creation of container images that can be used on a personal laptop or on a cloud platform. The workflow container image generated is freely available from the Google Container registry (<https://cloud.google.com/container-registry>) with name: *gcr.io/cbb-research-dl/transannot-cloud-cmp*

All files used to generate this image are available at: <https://github.com/ncbi/cloud-transcriptome-annotation/tree/master/config/gcp/docker>

Common Workflow Language

Common Workflow Language (CWL) [21] is an open standard workflow language used to describe and implement complex pipelines which uses interchangeable blocks. The resulting product is portable and scalable. It can be executed across a variety of hardware environments as dissimilar as personal laptops or the cloud.

Workflow managers are tools that simplify the execution of workflows in multiple computational environments. Some have been developed to manage and execute CWL workflows like Toil[23], CWL-Airflow [24], Arvados (<https://arvados.org/>) and REANA (<http://reanahub.io/>). Others, however, use their own workflow languages like Nextflow [25], and SnakePipes [26]. All provide a unified interface to users to choose the compute environment to process jobs. Users can configure the workflow manager to submit jobs to a high-performance compute cluster or to a cloud provider. Nevertheless, all these workflow managers use the cloud batch system to submit jobs for computing in the cloud.

In this study, we aim to estimate the minimum cost of executing a transcriptome annotation pipeline in the cloud. We selected CWL because it is the workflow language with many available workflow managers. Also, CWL provides a reference implementation runner: *cwltool* (<https://github.com/common-workflow-language/cwltool>). This runner can be executed on the command line inside a GCP or AWS job definition minimizing all dependencies for processing a workflow. We intentionally avoided the use of workflow managers to be able to quantify runtime for the workflow steps as precise as possible.

GCP

The Google Cloud Platform (GCP) offers a batch system specifically designed for life sciences, the Cloud Life Sciences (<https://cloud.google.com/life-sciences>). This system was initially Google Genomics but has evolved to allow the scientific community to process biomedical data at scale.

Cloud Life Sciences offers an Application Program Interface (API) implemented for users to develop their own workflow in JSON format using three main attributes: *actions*, *environments* and *resources*. *Actions* are the list of *commands* to execute using a defined container image. They also include statements to mount local solid-state drives (SSD) or network storage devices, defined in *resources*. *Environments* define the environment variables available inside the container. Finally, *resources* define the instance type and the local SSD or network storage devices.

The API, using the JSON described in Box 1, automatically creates instances on-demand, following the requirements defined in the *resources* section of the job JSON file. GCP also provides a customized container image where the instance interacts with other GCP products like Google Storage where data is stored. In addition, GCP creates the instances using a customized Linux operating system that formats and mounts the instance local disks making them available for the jobs.

Box 1 shows a brief extract of the pipeline used in GCP. We show only the main activity where the *command* attribute defines the command line to execute the CWL workflow. *ImageUri* attribute defines the container image used to run the command. In this case, our previously created Docker image. Finally, the *mounts* attribute defines the paths in the container to mount the disks created in the *resources* attribute.

Box 1: Brief extract of the GCP pipeline definition JSON file

```
{
  "actions": [
    ...,
    {
      "commands": [
        "/bin/bash",
        "-c",
        "cwltool --no-container --on-error continue --tmpdir-prefix /data/ --tmp-outdir-prefix /data/ --outdir /data/${SAMPLE}
https://raw.githubusercontent.com/ncbi/cloud-transcriptome-annotation/master/bin/cwl-ngs-workflows-
cbb/workflows/Annotation/transcriptome_annotation.cwl

--blast_db_dir /data --threads ${CPUs} --eval 1e-5 --blast_nt_db nt --blast_nr_db nr --blast_cdd_db split-cdd --fasta
/data/${SAMPLE}.fa >> /data/pipeline.log 2>&1"
      ],
      "imageUri": "gcr.io/cbb-research-dl/transannot-cloud-cmp",
      "mounts": [
        {
          "disk": "gcloud-shared",
          "path": "/data"
        }
      ]
    },
    ...,
    "environment": {
      "CPUs": "64"
    },
    "resources": {
      "virtualMachine": {
        "bootDiskSizeGb": 60,
        "bootImage": "projects/cos-cloud/global/images/family/cos-stable",
        "disks": [
          {
            "name": "gcloud-shared",
            "sizeGb": 600,
            "type": "local-ssd"
          }
        ],
        "machineType": "n1-standard-64",
        ...,
      }
    }
  ]
}
```

The *VirtualMachine* attribute defines the resources used to create the job instance. In this attribute, users can define instance boot disk size, operating system, extra disks and the machine type. The complete JSON file is available at: <https://github.com/ncbi/cloud-transcriptome-annotation/blob/master/config/gcp/pipeline.json>

AWS

AWS Batch (<https://aws.amazon.com/batch/>) is the batch system provided by Amazon Web Services. It is comprised of *compute environments*, *job queues* and *job definitions*. The *compute environment* defines the computational resources to be used by the batch jobs. It is connected to the Amazon Elastic Container Service (ECS) which is a fully managed service that creates and manages computer clusters inside the Amazon cloud environment. The resources defined by the compute environment are used by the ECS to create and setup instances in which the workload is distributed. *Job queues* are used as an intermediate service to associate submitted jobs with the compute environments. Lastly, the jobs use a *job definition*, in JSON format, which defines specific information for the job, like container images, commands, number of vCPUs, RAM memory, environment variables and local or remote folder to mount on the container.

Box 2 shows a brief extract of the job definition JSON script used in AWS. The *containerProperties* attribute defines the job properties. *Image* defines the container image, in this case our Docker image. *Command* defines the command to be executed inside the container. In the case of AWS, a single command can be outlined in the job definition, thus, complex pipelines with multiple steps can be encapsulated in a BASH script. This script can be stored inside the container image or the container can download it at runtime. For simplicity, we have included this script inside the Docker image.

Box 2: Brief extract of the AWS job definition JSON file

```
{
  ...,
  "containerProperties": {
    "image": "gcr.io/cbb-research-dl/transannot-cloud-cmp",
    "vcpus": 64,
    "memory": 131072,
    "command": [
      "/usr/envs/transannot/bin/aws-pipeline.sh"
    ],
    "volumes": [
      {
        "host": {
          "sourcePath": "/data"
        },
        "name": "data"
      }
    ],
    "environment": [
      {
        "name": "CPUs",
        "value": "32"
      }
    ],
    "mountPoints": [
      {
        "containerPath": "/data",
        "sourceVolume": "data"
      }
    ],
    ...,
  }
}
```


The AWS Batch system automatically creates all infrastructure, network components and compute instances, following the requirements of the compute environments. The default configuration of the Amazon Machine Image (AMI) used for the instances, however, is not configured to use local SSD disks available on certain machine types. This limits the default options on the AWS Batch system to certain types of workflows. Workflows that use intensive disk IO operations will have improved performance and efficiency if local SSD disks are used. Thus, a modified AMI capable of use the instance local disks is required for our study. We create a customized AMI for our study that is freely available in the AWS zone *us-east1* with ID: *ami-0dac0383cac1dc96e*. This AMI creates an array with the local SSD disks in the instance using the Linux utility *mdadm*. The array is formatted with XFS filesystem and mounted in a folder named */data*.

To improve the default AWS Batch options, Amazon offers a Virtual Private Cloud (VPC) that provides an extra layer of isolation for the resources used by the AWS Batch system. This VPC logically isolates all resources used in a defined virtual network improving the security. It is customizable for each compute problem.

The templates used in our study to create all the components of the AWS Batch system are available at: <https://github.com/ncbi/cloud-transcriptome-annotation/tree/master/config/aws>. All resources are created in the Jupyter notebook in: “02 - AWS-Batch”.

GCP and AWS batch system limitations

Bioinformatics best practices for pipeline execution require the containerization of each tool included in the analysis. Projects such as Bioconda [27] and Biocontainers [28] provide standard containerized images for thousands of bioinformatics tools. However, the batch system for both tested cloud providers requires that all tools used in the workflow to be included in a single container. Each action in the cloud job definition has associated a single Docker image that is used to execute the action task. Docker-in-Docker, the process to execute docker containers inside another Docker container is not permitted in both GCP and AWS. This limitation constrains users to containerize all tools involved in a workflow, into a single Docker image. Hence, knowledge on how to create Docker images is a requirement for the migration of workflows to the cloud.

GCP and AWS transitory instances

Both GCP and AWS offer access to transitory instances which are spare compute capacity at a reduced cost. These instances are called SPOT in AWS and Preemptible in GCP. The transitory instances at reduced cost results from the fact that the cloud provider might terminate the instance at any time. Preemptible prices in GCP are fixed but not in AWS. The cost of the SPOT instances has a minimum but can be increased to the normal EC2 price if the demand for resources increases.

Transitory instances for workflow execution require extra processing steps to identify terminated jobs for resubmission. This is a reasonable option to reduce the cost of the analysis but requires a flexible timeframe to complete all analyses. Users need to be aware of this caveat.

Jupyter Notebooks

Jupyter notebooks are an open-source web application framework for the creation and sharing of documents that contain live code [29]. It is a standard way to share scientific code for ease of reproducibility and reuse [30]. The implementation of our study was fully developed in Jupyter notebooks. Readers can reproduce our results and figures using the notebooks that are available at the project GitHub repository. The notebooks create all cloud resources and submit the jobs to the batch systems. They also retrieve the job logs in JSON format and create the figures automatically from those logs. Each notebook includes a description about its purpose and is named using a numeric prefix to highlight the execution order.

The notebooks implemented in this study are designed to be executed on a local laptop or a workstation. Both interact asynchronously with the cloud providers using the command line APIs provided. In the case of GCP, we used the Google Cloud SDK (<https://cloud.google.com/sdk>). For AWS, we used the AWS Command Line Interface (<https://aws.amazon.com/cli/>). In these notebooks, the workflow input files are created, these are uploaded to each cloud provider storage space, the cloud batch systems are configured, the jobs are submitted and the results are retrieved. The notebooks interact with the cloud batch system to process jobs and retrieve results and logs stored in the *results/PRJNA320545* folder.

Table 1: Machine types with resources in each cloud

Provider	Machine type	vCPU	Memory (GB)	Instance Local SSD (GB)	Network Bandwidth (Gbps)	USD/Hour
AWS	m5d	16	64	2 x 300	Up to 10	0.904
AWS	m5d	32	128	2 x 600	10	1.808
AWS	m5d	64	256	4 x 600	20	3.616
AWS	m5dn	16	64	2 x 300	Up to 25	1.088
AWS	m5dn	32	128	2 x 600	25	2.176
AWS	m5dn	64	256	4 x 600	75	4.352
GCP	n1	16	60	24 x 375	32	0.861
GCP	n1	32	120	24 x 375	32	1.393
GCP	n1	64	240	24 x 375	32	2.475
GCP	n2	16	64	24 x 375	32	0.951
GCP	n2	32	128	24 x 375	32	1.572
GCP	n2	64	256	24 x 375	32	2.816

Results and Discussion

In this study, we present an analysis of the complexity, cost and best practices for executing the core components of a transcriptome annotation workflow in the cloud. For our experiments, we used the two cloud provider partners of the NIH’s STRIDES Initiative: GCP and AWS. For each cloud provider, similar compute instances were tested using 16, 32 and 64 vCPUs. The machine types and their resources are described in Table 1. We used the transcriptome assembled from a public BioProject with ID PRJNA320545 for the organism *Opuntia streptacantha* (prickly pear

cactus). The transcriptome includes 474,563 transcripts generated with Trinity [31], and is available in [data/PRJNA320545/transcriptome.fasta.gz](https://data.prjna320545/transcriptome.fasta.gz). The transcriptome length distribution and statistical metrics are available in the 01 - Data Partitioning notebook.

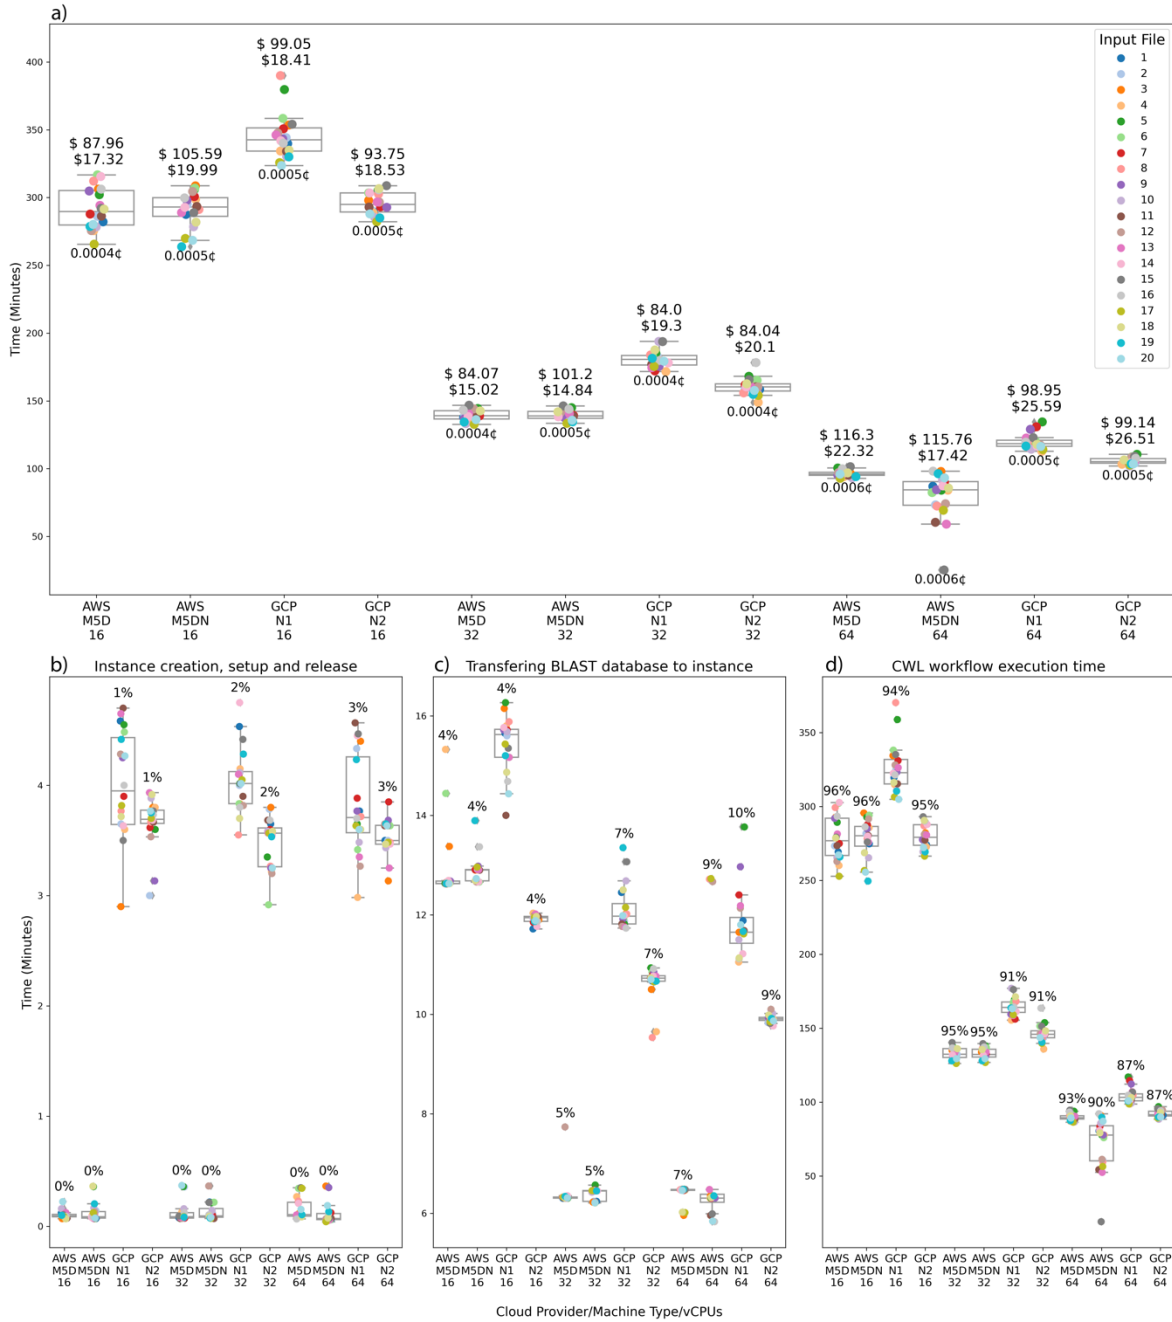


Figure 3: Time and cost for the 10,000 query size files. a) Total time for each input file for each configuration (Cloud provider/Machine Type/vCPUs). The total cost of processing the 20 input files (200,000 transcripts in total) is at the top of each box using normal and transitory instances. The cost of processing one transcript is at the bottom of each box. b) Time and percent of the total cost for instance creation, setup and release. c) Time and percent of the cost for transferring the BLAST databases to the instance from the cloud storage bucket (S3 in AWS and Cloud Storage in GCP). d) Time and percent of the cost for the CWL workflow execution.

From the *Opuntia* pool of transcripts, we analyzed three sizes of *query* files: 2,000, 6,000, and 10,000 transcripts in each input *query* file. Two experiments were executed. First 20 FASTA files (input files for the workflow) for each query size were randomly created, see notebook “01 - Data Partitioning”. Each of these files were submitted independently as jobs to the batch systems on each cloud provider. For the second experiment, 120,000 transcripts were randomly selected and then partitioned in files with 2,000, 6,000, and 10,000 transcripts to analyze the relationship between query size, runtime and cost.

Jobs were submitted to each cloud platform using the notebook “02 - Google Cloud Platform” and “02 - AWS-Batch”. In each notebook, the input files created for each experiment were copied to the respective cloud storage system, followed by job submissions for each configuration of machine type/CPU.

Four times were collected from the jobs:

1. the total run time
2. the time to transfer the BLAST databases to the instance local SSD disk
3. the time executing the CWL workflow
4. the time for creation, setup and release of the instance

Figure 3 shows the collected times for the 10,000-query size. In Figure 3a the total run time for each input file (each containing 10,000 transcripts) for a total of 200,000 transcripts processed for each cloud provider, machine type and the number of vCPUs.

In addition, each box in Figure 3a shows the total cost for the 20 files using normal and transitory instances (top) and the cost of processing one transcript (bottom). The bottom row with three plots shows the remaining three times collected from the jobs.

The total running time for the 10,000-query sized files are similar for the same number of vCPUs notwithstanding the cloud provider. Furthermore, this example shows how the running time can be reduced by more than a half by increasing the number of vCPUs. Unfortunately, this time reduction does not decrease the total cost of the project as the price per hour for machines with more vCPUs increases as well.

The AWS platform is more efficient than the GCP during the instance creation, setup and release, see Figure 3b. This stage takes only 0.1% of the total cost. The GCP cost for this stage goes from 1.5% to 4.5% on bigger machines. The differences are due to the Amazon Elastic Container Service (ECS) which allocates new jobs on existing instances as soon as the instance gets free without releasing them, whereas GCP creates, sets up, and releases an instance for each job.

Transferring the BLAST databases from each cloud storage (S3 in AWS and Cloud Storage in GCP) bucket, Figure 3c, (current size is 342GB), to the instance local SSD disk is a crucial step in reducing the cost of the analysis. Initially, we tested the default parameters in both cloud providers which use network storage devices taking an average of 1 hour which is about 30 % of the total cost of the analysis and takes more time than the CWL workflow execution. After customizing both batch systems to use the instance local SSD disks, the time was reduced to a range of 4 % to 11 % of the total cost in the 10,000 query size.

As expected, the CWL workflow execution time is the most time-consuming part of the job, Figure 3d. All configurations show similar times for executing the CWL workflow. The GCP N1

machine type spent more time on the CWL workflow than the other machine types in all configurations because the GCP N1 is the Google first generation machine type with slower vCPUs.

Figure 4 shows the time and cost of processing 120,000 transcripts using second generation 64 vCPUs instances on each cloud provider. Reducing the number of transcripts per input file reduces the total run time but will also increase the cost of the analysis as more instances will be used. BLAST databases are transferred to more instances spending, on average, 10 minutes for each instance. For example, our experiment with the 10,000-query size processes all transcripts in about 105 minutes with a total cost of 59.37USD using 12 instances (GCP, N2, 64 CPU). Processing the same number of transcripts with a query size of 2,000 costs 122.36USD with all transcripts processed in 43 minutes using 60 instances (GCP, N2, 64 CPU).

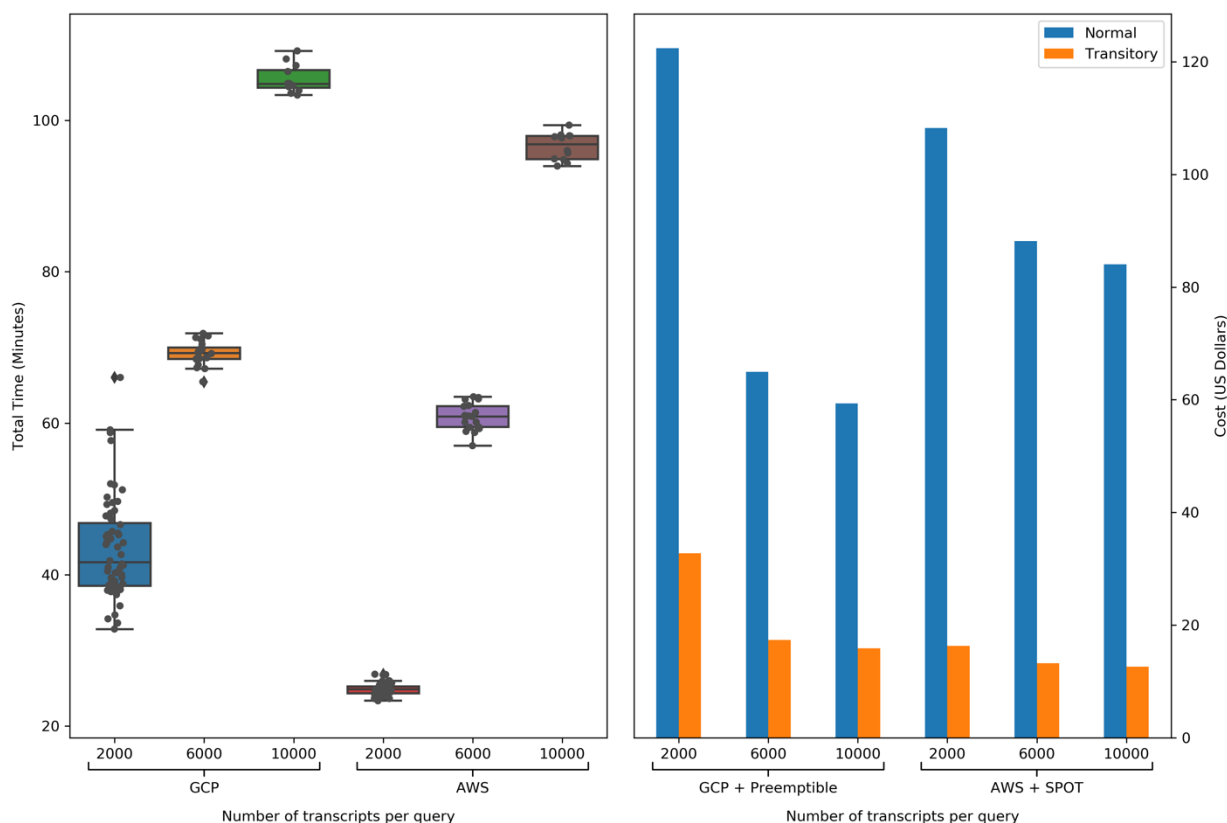


Figure 4: Left plot shows the total processing time for 120,000 transcripts using different query sizes. Right plot shows the total cost using normal compared to transitory instances.

We have determined that a transcriptome with ~500,000 transcripts can be processed in less than 2 hours with a compute cost ranging from 200USD to 250USD using normal instances. For transitory instances (SPOT in AWS and preemptible in GCP) the total cost could be reduced to 50USD for the complete analysis. However, the processing of all transcripts requires a flexible timeframe due to the availability of the transitory instances and the number of terminated jobs that require resubmission. In our opinion, these are reasonable costs that make the transcriptome annotation process in the cloud accessible to any genomic laboratory without access to an on-premise computational infrastructure.

Best practices

Our recommendation for best practices using public cloud providers for computational biology experiments are:

1. For reproducibility, write the pipeline using a workflow language.
We recommend CWL because the resulting product is portable and scalable, and it can be executed across a variety of computational environments as dissimilar as personal laptops or the cloud. As mentioned above, CWL is the workflow language with many workflow managers available and they can be directly executed in a container using the **cwltool** runner.
2. Containerize the CWL workflow with Docker.
Use Conda/Bioconda to install all Bioinformatics tools in the container image.
3. Use Jupyter Notebooks for coding and documenting each step during experiments.
4. Use the cloud provider batch system for deploying jobs.
5. Cloud computing behaves differently than local workstations or on-premise clusters.
Users should define and execute small tests with their data and workflow before submitting large jobs. Testing different cloud services and configurations could help to reduce the runtime and cost for the whole analysis.
6. Use the instance local disks for computing instead the default network devices.
7. Use transitory instances to reduce the cost **ONLY** if there are no timeframe restrictions for completing the analysis.

Conclusion

Despite differences in the configuration and setup of batch systems between GCP and AWS, the cost and processing time are similar for the type of workflow we designed for our experiment. In our opinion, for BLAST-based workflows, the choice of a cloud platform is not dependent on the workflow but, rather, on the specific details of the cloud provider. These specific details are related to the accessibility of each cloud platform for institutional use, the technical knowledge of the specific platform services, and/or the availability of open-source frameworks to deploy the workflows on a specific cloud provider.

We found that GCP is easier to use as it only requires a JSON file for batch processing whereas AWS needs a complete setup of all batch system components. GCP is more suitable for daily data analysis work in research laboratories. On the other hand, AWS, once properly configured, is more efficient in terms of machine creation, setup and release. The ECS can reuse instances reducing the cost for large data analysis projects. AWS is more suitable for large data analysis groups to establish a set of queues and compute environments for multiple pipelines.

Availability of supporting source code and requirements

Project name: Cloud Transcriptome Annotation

Project home page: <https://github.com/ncbi/cloud-transcriptome-annotation>

Operating system(s): Linux and MacOS

Programming languages: Python, BASH

Other requirements: Conda/Bioconda, Jupyter Notebook

CWL workflow:

https://github.com/ncbi/cloud-transcriptome-annotation/blob/master/bin/cwl-ngs-workflows-cbb/workflows/Annotation/transcriptome_annotation.cwl

CWL Viewer:

https://w3id.org/cwl/view/git/0d8650062673c8af2c1139c557afc4c3d6a1b53c/bin/cwl-ngs-workflows-cbb/workflows/Annotation/transcriptome_annotation.cwl

Abbreviations

AMI: Amazon Machine Images (AMI)

API: Application Program Interface

AWS: Amazon Web Services

CDD: Conserved Domain Database

CWL: Common Workflow Language

ECS: Amazon Elastic Container Service

GCP: Google Cloud Platform

NCBI: National Center for Biotechnology Information

ORFs: Open Reading Frames

SSD: Solid State Disk

VPC: Virtual Private Cloud

Competing interests

The authors declare that they have no competing interests.

Authors' contributions

RVA, LMR and DL contributed to the design of the annotation workflow and the manuscript preparation. RVA designed, implemented and executed all cloud environments, configurations and experiments. All authors read and approved all versions of the manuscript.

Funding

This work was supported by the Intramural Research Program of the National Library of Medicine, National Center for Biotechnology Information at the National Institutes of Health.

Acknowledgements

We would like to thank:

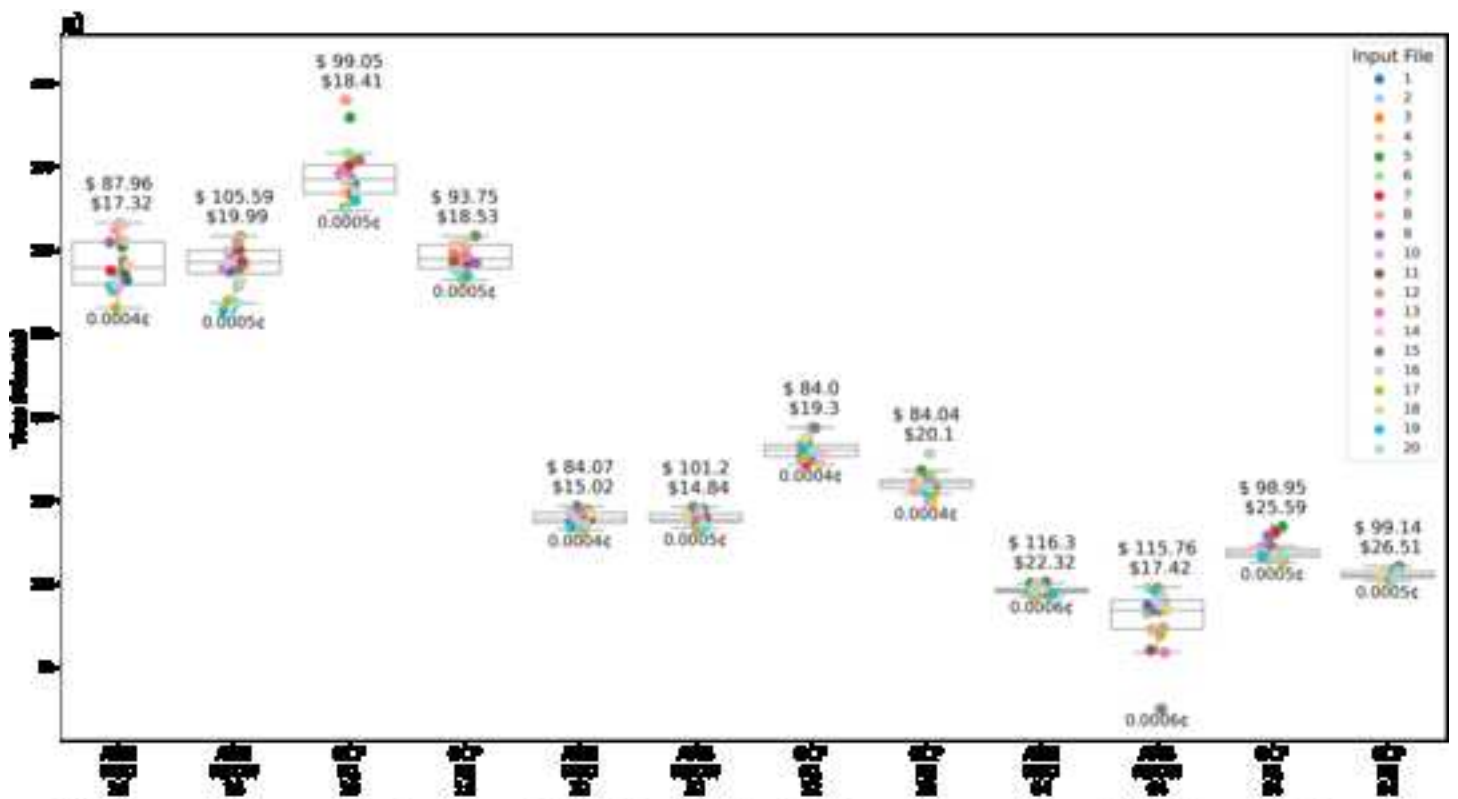
NCBI BLAST Group: Christiam Camacho, Vadim Zalunin, Greg Boratyn, Ryan Connor and Tom Madden for their support with BLAST.

NCBI Cloud and System Group: Al Graeff, Brian Koser, Andrew Arensburger, Brad Plecs, Ron Patterson and Dima Beloslyudtsev for their support with the cloud platforms.

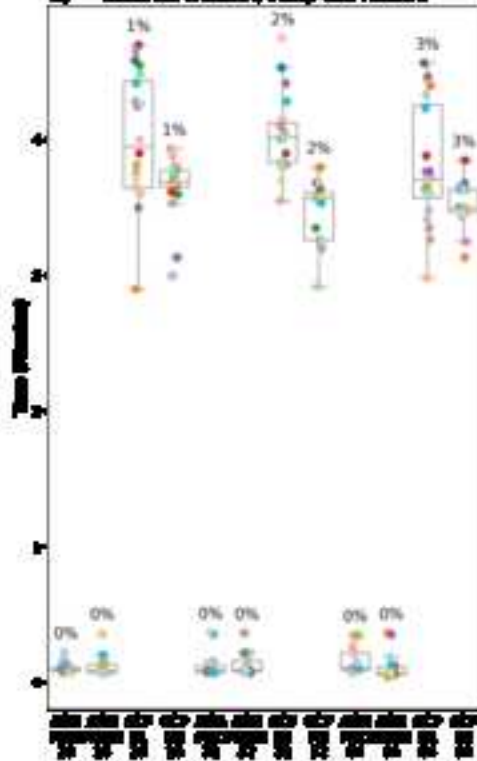
References

1. Langmead, B. and A. Nellore, *Cloud computing for genomic data analysis and collaboration*. Nature Reviews Genetics, 2018. **19**(4): p. 208-219.
2. Sayers, E.W., et al., *Database resources of the National Center for Biotechnology Information*. Nucleic Acids Res, 2020. **48**(D1): p. D9-D16.
3. Al-Qurainy, F., et al., *Comprehensive Stress-Based De Novo Transcriptome Assembly and Annotation of Guar (Cyamopsis tetragonoloba (L.) Taub.): An Important Industrial and Forage Crop*. Int J Genomics, 2019. **2019**: p. 7295859.
4. Chabikwa, T.G., et al., *De novo transcriptome assembly and annotation for gene discovery in avocado, macadamia and mango*. Sci Data, 2020. **7**(1): p. 9.
5. Ji, P., et al., *Characterization of common carp transcriptome: sequencing, de novo assembly, annotation and comparative genomics*. PLoS One, 2012. **7**(4): p. e35152.
6. Torre, S., et al., *RNA-seq analysis of Quercus pubescens Leaves: de novo transcriptome assembly, annotation and functional markers development*. PLoS One, 2014. **9**(11): p. e112487.
7. Carruthers, M., et al., *De novo transcriptome assembly, annotation and comparison of four ecological and evolutionary model salmonid fish species*. BMC Genomics, 2018. **19**(1): p. 32.
8. Haas, B.J., et al., *De novo transcript sequence reconstruction from RNA-seq using the Trinity platform for reference generation and analysis*. Nat Protoc, 2013. **8**(8): p. 1494-512.
9. Bryant, D.M., et al., *A Tissue-Mapped Axolotl De Novo Transcriptome Enables Identification of Limb Regeneration Factors*. Cell Rep, 2017. **18**(3): p. 762-776.
10. Vera Alvarez, R., et al., *Workflow and web application for annotating NCBI BioProject transcriptome data*. Database (Oxford), 2017. **2017**.
11. Gamez, R.M., et al., *Banana (Musa acuminata) transcriptome profiling in response to rhizobacteria: Bacillus amyloliquefaciens Bs006 and Pseudomonas fluorescens Ps006*. BMC Genomics, 2019. **20**(1): p. 378.
12. Altschul, S.F., et al., *Basic local alignment search tool*. J Mol Biol, 1990. **215**(3): p. 403-10.
13. Ashburner, M., et al., *Gene ontology: tool for the unification of biology. The Gene Ontology Consortium*. Nat Genet, 2000. **25**(1): p. 25-9.
14. Peters, K., et al., *PhenoMeNal: processing and analysis of metabolomics data in the cloud*. Gigascience, 2019. **8**(2).
15. Belyeu, J.R., et al., *SV-plaudit: A cloud-based framework for manually curating thousands of structural variants*. Gigascience, 2018. **7**(7).
16. Kiar, G., et al., *Science in the cloud (SIC): A use case in MRI connectomics*. Gigascience, 2017. **6**(5): p. 1-10.
17. Hiltmann, S., et al., *CGtag: complete genomics toolkit and annotation in a cloud-based Galaxy*. Gigascience, 2014. **3**(1): p. 1.

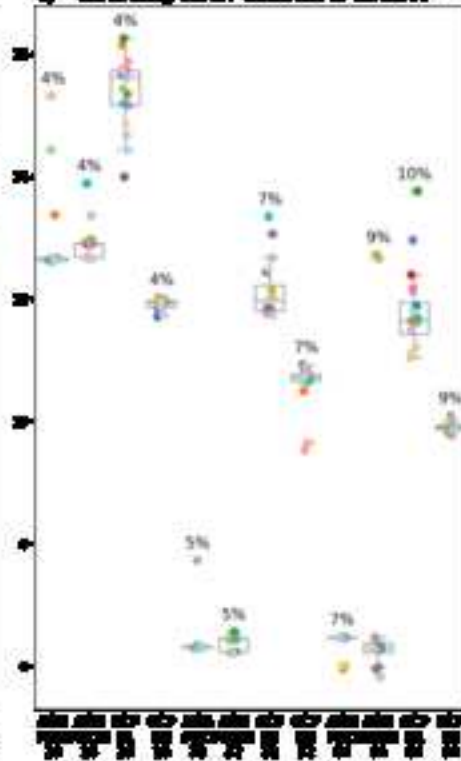
18. Ohta, T., T. Tanjo, and O. Ogasawara, *Accumulating computational resource usage of genomic data analysis workflow to optimize cloud computing instance selection*. *Gigascience*, 2019. **8**(4).
19. Haas, B. and A. Papanicolaou. *TransDecoder (Find Coding Regions Within Transcripts)*. 2020; Available from: <https://github.com/TransDecoder/TransDecoder/wiki>.
20. Yang, M., et al., *NCBI's Conserved Domain Database and Tools for Protein Domain Analysis*. *Curr Protoc Bioinformatics*, 2020. **69**(1): p. e90.
21. Peter, A., et al., *Common Workflow Language, v1.0*. 2016.
22. Pertea, M., *The human transcriptome: an unfinished story*. *Genes (Basel)*, 2012. **3**(3): p. 344-60.
23. Vivian, J., et al., *Toil enables reproducible, open source, big biomedical data analyses*. *Nat Biotechnol*, 2017. **35**(4): p. 314-316.
24. Kotliar, M., A.V. Kartashov, and A. Barski, *CWL-Airflow: a lightweight pipeline manager supporting Common Workflow Language*. *Gigascience*, 2019. **8**(7).
25. Di Tommaso, P., et al., *Nextflow enables reproducible computational workflows*. *Nat Biotechnol*, 2017. **35**(4): p. 316-319.
26. Bhardwaj, V., et al., *snakePipes: facilitating flexible, scalable and integrative epigenomic analysis*. *Bioinformatics*, 2019.
27. Gruning, B., et al., *Bioconda: sustainable and comprehensive software distribution for the life sciences*. *Nat Methods*, 2018. **15**(7): p. 475-476.
28. da Veiga Leprevost, F., et al., *BioContainers: an open-source and community-driven framework for software standardization*. *Bioinformatics*, 2017. **33**(16): p. 2580-2582.
29. Shen, H., *Interactive notebooks: Sharing the code*. *Nature*, 2014. **515**(7525): p. 151-2.
30. Perkel, J.M., *Why Jupyter is data scientists' computational notebook of choice*. *Nature*, 2018. **563**(7729): p. 145-146.
31. Grabherr, M.G., et al., *Full-length transcriptome assembly from RNA-Seq data without a reference genome*. *Nat Biotechnol*, 2011. **29**(7): p. 644-52.



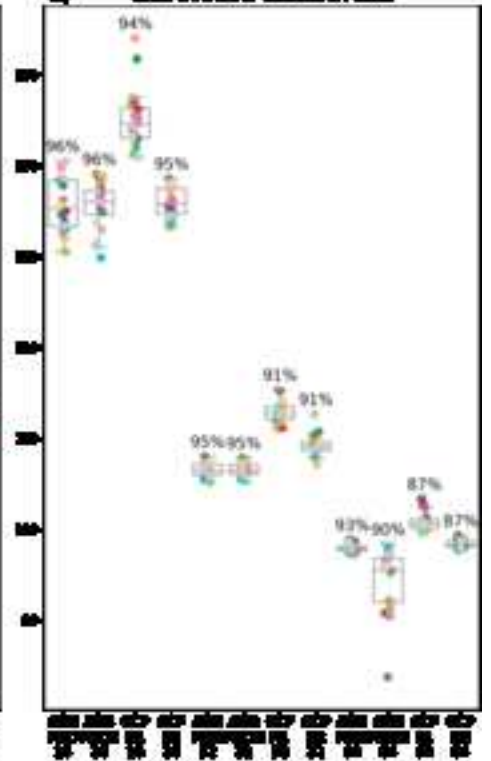
b) Instance creation, setup and release



c) Transferring ELASTIC database to instance

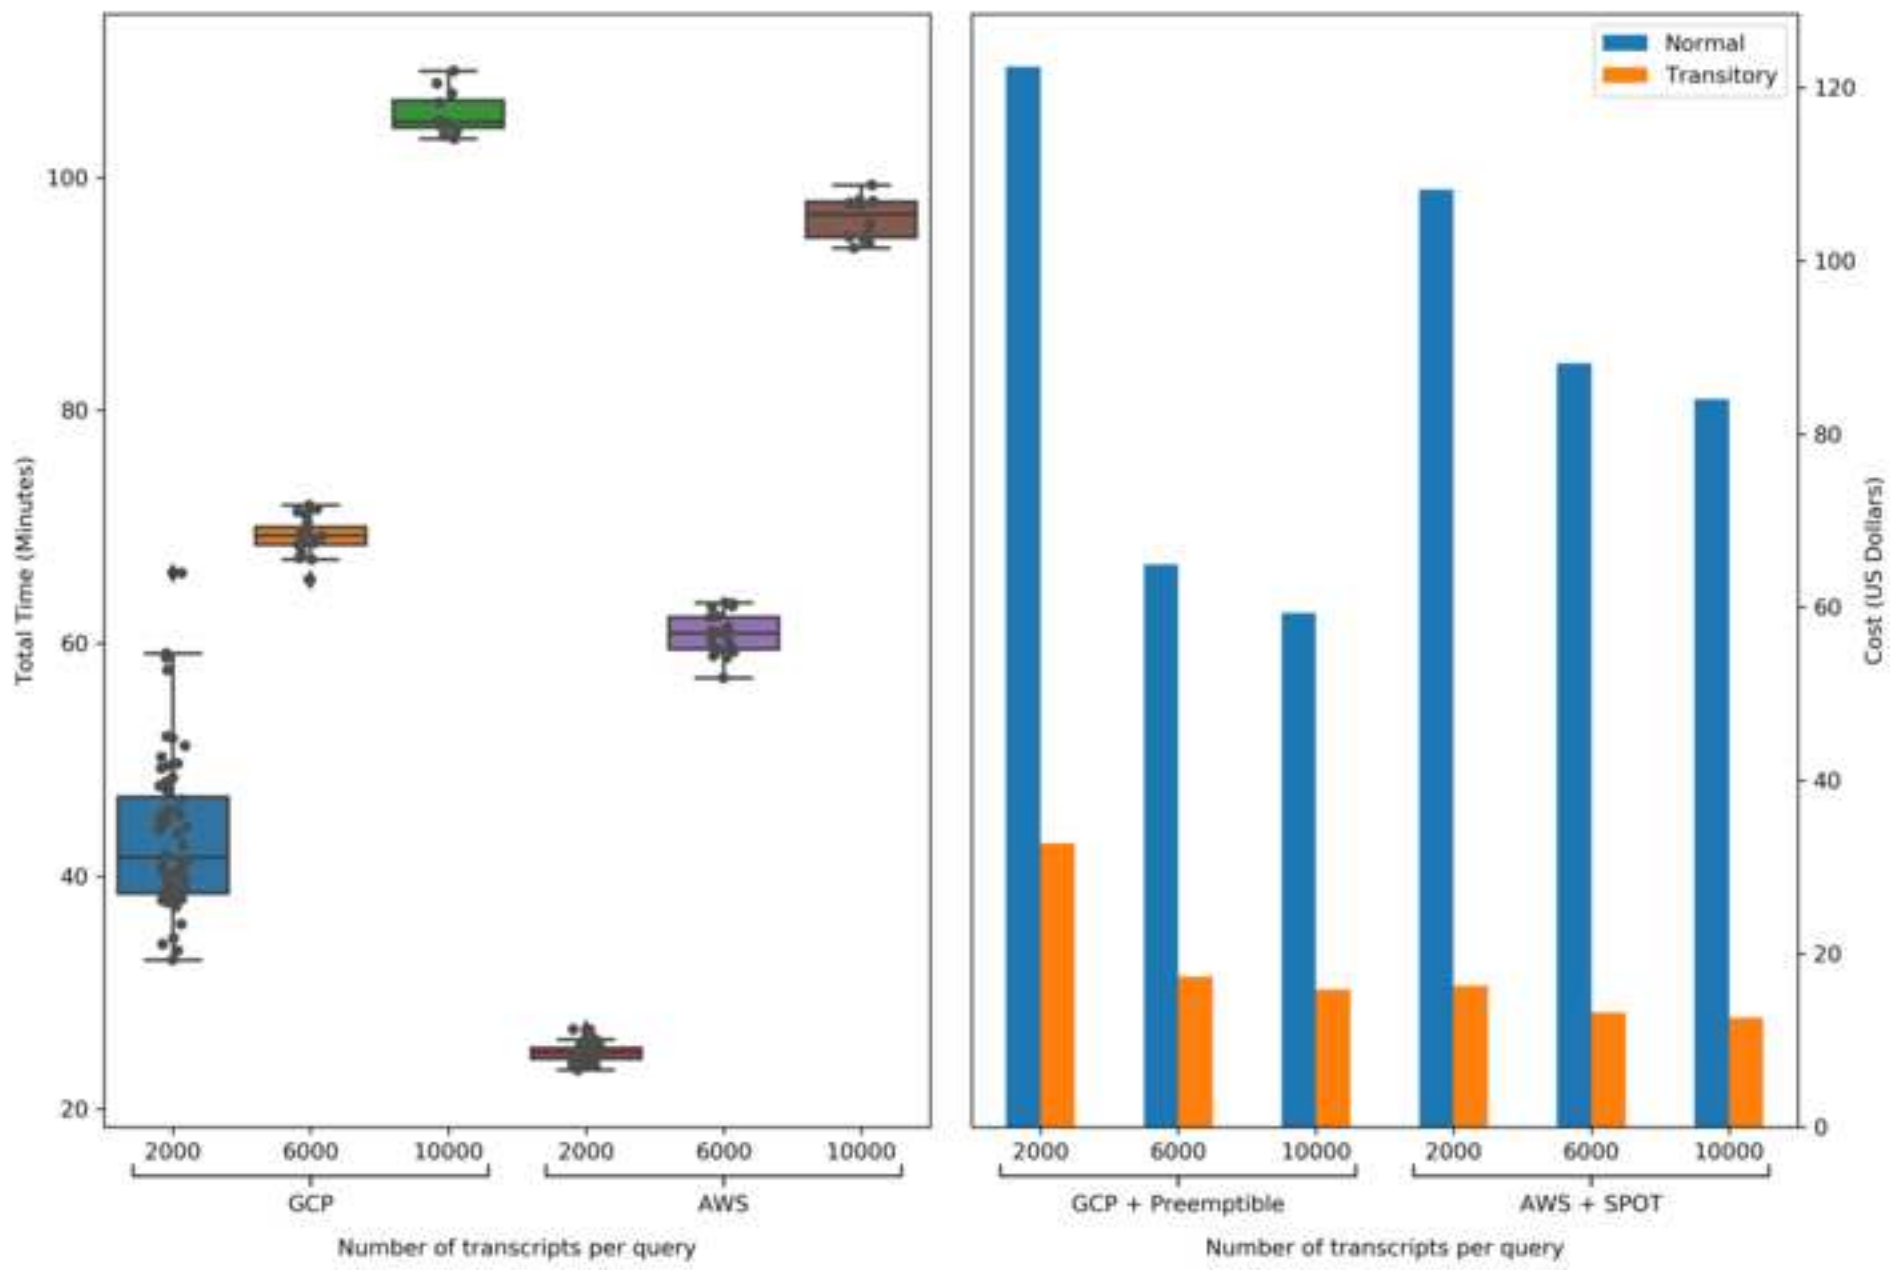


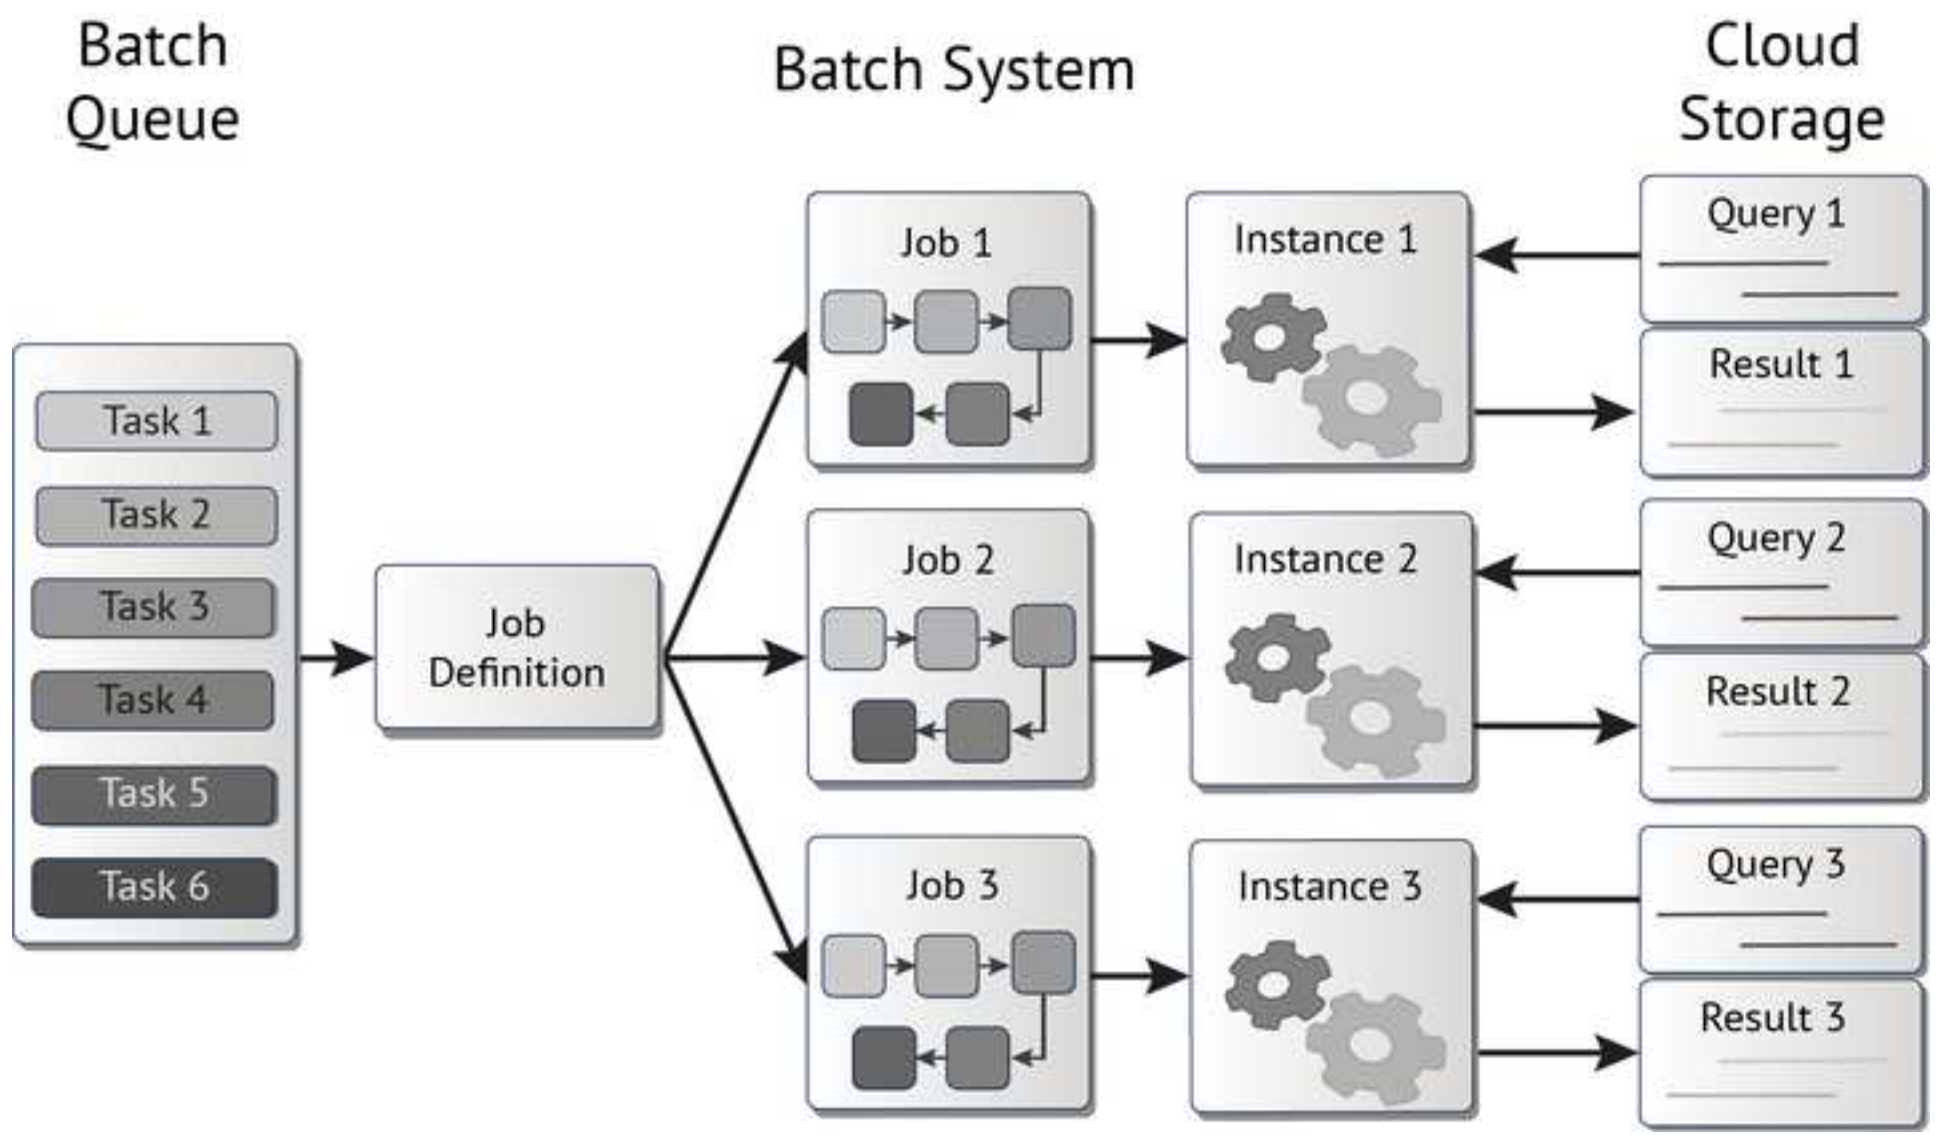
d) CML workflow execution time

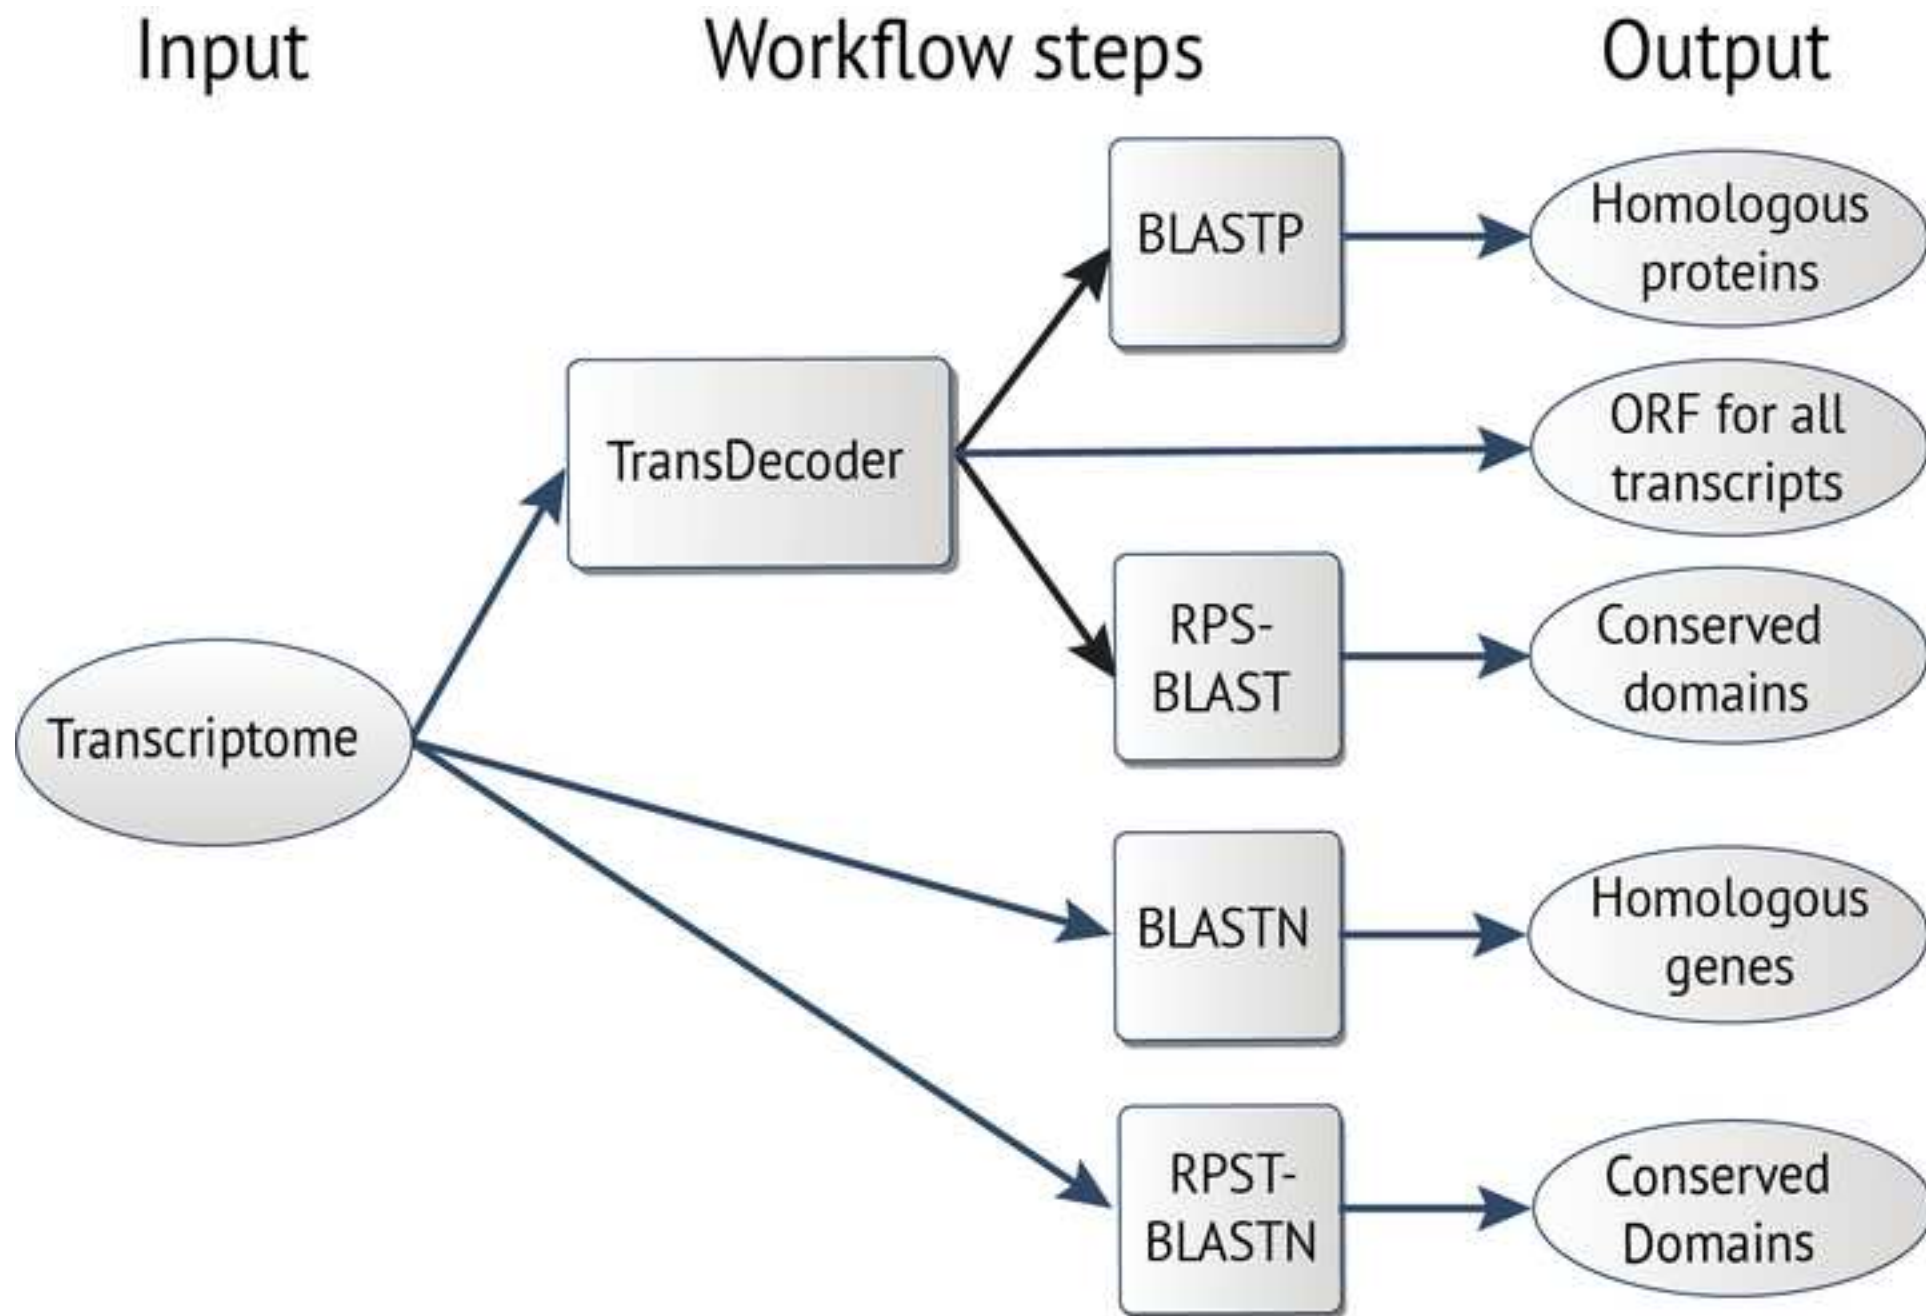


Cloud Provisioning Time/CPU

Figure 4









DEPARTMENT OF HEALTH & HUMAN SERVICES

Public Health Service

National Institutes of Health
National Library of Medicine
Bethesda, Maryland 20894

November 13, 2020

Dear Editors,

We would like to resubmit our manuscript entitled “Transcriptome annotation in the cloud: complexity, best practices and cost” for your consideration for publication in *GigaScience as a Technical Note*.

We thank the editors and referees for the considerable time dedicated to our manuscript. The referees’ revisions were particularly constructive and have helped us to improve our study. We think that we have answered **all** their questions and addressed **all** their comments in the manuscript.

In this manuscript, we present a comparative study of multiple BLAST alignments using two public cloud providers: Amazon Web Services (AWS; Seattle, WA, USA) and Google Cloud Platform (GCP; Mountain View, CA, USA). These are the cloud providers partners of the **NIH Science and Technology Research Infrastructure for Discovery, Experimentation, and Sustainability (STRIDES) initiative**. We have prepared several Jupyter Notebooks with all the code required to submit BLAST jobs to the batch system on each cloud provider in order to reproduce or extend our results. We demonstrate that the public cloud providers are a practical alternative for executing advanced computational biology experiments at quite low cost. Using our cloud recipes, the BLAST alignments required to annotate a transcriptome with ~500,000 transcripts can be processed in less than 2 hours with a computing cost of about US\$ 200-250.

All code used during this study is freely available at: <https://github.com/ncbi/cloud-transcriptome-annotation>

Thank you very much for your consideration of our manuscript. We look forward to your reply on this submission.

Sincerely,

Roberto Vera Alvarez ^{1,a}, Leonardo Mariño-Ramírez ^{1,2,b} and David Landsman ^{1,c}

¹ Computational Biology Branch, National Center for Biotechnology Information, National Library of Medicine, National Institutes of Health, Bethesda, MD, USA.

² Current address: Division of Intramural Research, National Institute on Minority Health and Health Disparities, NIH, Bethesda, MD, USA.

- a. veraalva@ncbi.nlm.nih.gov
 - b. marino@nih.gov
 - c. landsman@ncbi.nlm.nih.gov
-
-

Reviewer reports:

Response: We thank the referee for the comments and for their time dedicated to this manuscript.

Reviewer #1: In this manuscript, the authors describe the usability and the performance of public cloud computing services and their batch job execution services for the transcriptome annotation pipeline. The authors also compare the services of two major public cloud vendors, Google Cloud Platform (GCP), and Amazon Web Service (AWS). The results show that the two cloud providers can provide a similar experience on its operation, workflow execution time, and payment for execution. The performed experiments follow the modern best practice of data analysis using tools and frameworks for better reproducibility, including Docker container, the Common Workflow Language (CWL), and Jupyter Notebook. This article can be a good example of a reproducible study with open data and open-source software.

This report is very significant with the practical statistics, which can be a helpful reference for all the cloud use cases in biomedical data analysis. The title states that this study focuses on the transcriptome annotation, but the output provides insight for all the cloud use cases. I suggest the authors change the title because the current one looks the best practice valid only for the transcriptome annotation in the wide variety of genomic data analysis.

Response: We thank the referee for their comment and suggestions. It is true that the results of this study provide insight for many computational biology workflows but our experiments were limited to the transcriptome annotation process, specifically to the BLAST searches that are the core of the annotation. We think that the best practices and conclusions of this study should remain limited to the cloud transcriptome annotation process.

Below are minor comments for the manuscript.

Page 3, the last paragraph:

The authors claim as "little has been published describing cloud costs and implementation best practices". However, there is a study implemented software to monitor the runtime metrics of a given workflow, and support the cost estimation for the executions on the cloud (<https://doi.org/10.1093/gigascience/giz052>). This article describes the cost for the normal EC2 instance and does not mention the batch execution services, yet it provides additional information to the readers. Please consider introducing this in the background section as a related study. (Disclaimer: I am the first author of this article)

Response: We agree with the referee that this study provides additional information about the use and cost of computational biology workflows in the cloud. A new paragraph was added to the Background section: "The utilization of cloud environments for computational biology experiments is increasing [14-17], however, little has been published estimating cloud costs and implementation best practices. A recent work published by Ohta *at al.* [18] presents a tool

named CWL-metrics that collects runtime metrics of Docker containers and workflow metadata to analyze workflow resource requirements. This study presents a cost estimation for the execution on the cloud for AWS EC2 instances, but does not mention the cloud batch system for users to submit thousands of jobs to the cloud.”

Page 12, Best practices

The authors describe here that they recommend CWL as a workflow language. I think the authors should introduce the reason they chose CWL at the Method section where they first mention CWL (Page 4, Transcriptome Annotation Workflow). The authors also should provide a more practical reason to choose CWL because CWL is not only one framework that has portability and scalability. For example, having multiple workflow runners or its syntax easily parsed and connected to the resources such as runtime metrics can be a reason to choose CWL in this use case.

Response: We agree with the referee about the inclusion of more information about CWL and the reason why it was selected for our study. We added a new section to the manuscript named: **Common Workflow Language**

Workflow dependencies

Both AWS batch and Google LifeScience allow users to specify only one container per one batch job. This means the user needs to install the workflow runner (in this article cwltool) and the tools to process the data (e.g., BLAST) in a single container. However, the current best practice for bioinformatics is to separate the containers for each tool. Thus, in most cases, users cannot reuse the containers they use for a normal computing environment to the cloud batch system. The authors need to mention the limitation that one needs to create a container for cloud batch systems. Another possible solution would be to run sibling containers from the main batch job container, though I do not know if it is feasible with AWS and GCP.

Response: The referee’s comment is correct. Docker-in-Docker, the process to execute docker containers inside of another Docker container is not allowed in both GCP and AWS. We add a new section to the manuscript named **GCP and AWS batch system limitations** to address these limitations.

Software and framework related to the cloud batch services Many workflow languages and runners are supporting AWS batch and GCP. For example, Nextflow (<https://www.nextflow.io/docs/latest/awscloud.html>), Cromwell (<https://cromwell.readthedocs.io/en/stable/backends/Google/>), or Snakemake (<https://snakemake.readthedocs.io/en/stable/executing/cloud.html#executing-a-snakemake-workflow-via-tibanna-on-amazon-web-services>) can run the workflows via AWS batch or GCP. Task Execution Service (TES) of the Global Alliance for Genomics and Health (GA4GH) Cloud working group is also a framework to utilize the cloud batch system (<https://github.com/ga4gh/task-execution-schemas>). Some tools that run workflows on the cloud using the ETL framework (<https://doi.org/10.21105/joss.01069>) or cloud batch services

<https://github.com/DataBiosphere/dsub>) are also available. These may not be directly relevant to this study, but it would be helpful for readers to understand how the cloud batch services are being used by the researchers.

Response: We agree with the referee about the inclusion of more information about workflow managers. We added a new section to the manuscript named: **Common Workflow Language**.

Reviewer #2: Reviewer's report

Title: Transcriptome annotation in the cloud: complexity, best practices and cost. transcriptome data

Reviewers: Qiao Xuanyuan and Lucas B. Carey

Response: We thank the referees for their comments and time dedicated to this manuscript.

Reviewer's report:

The authors provide a perspective on transcriptome annotation in the cloud by presenting a comparative study of the two main public cloud providers, aiming to help the reader determine the proper cloud platform and its utilization in their research. The main strength of this paper is that it addresses a question that little has been studied before—the cloud cost estimates and implementation best practices. This is useful not only for labs doing transcriptome annotation, but, because all code is provided and very well commented, it might be of use for labs dealing with big data & distributed computation problems in general.

The manuscript is well written, and I have only a few minor concerns on presentation and the information that is provided.

1. The tested transcriptome data need to be clarified. I read the Data partitioning code for the creation of 20 FASTA files of the different query sizes (Fig 3). The variation in processing time among these files is consistent across clusters, and is therefore presumably due to differences in query sets in each file. This is surprising, as 10,000 is a large number. Are the differences because the sequence records were created sequentially from the input fasta file, with some genes having large numbers of hits? Would subsetting transcripts randomly result in more uniform processing times? There is almost two-fold variation in processing time between query files.

Response: The referee's comment is correct. The variation in processing time for queries with the same number of transcripts was due to the sequentially subsampling approach. We have modified the Data Partitioning notebook to create a random sampling of the transcriptome file. The notebook now shows the statistics for each file created. All measured parameters, mean, standard deviation, minimum length, 25%, 50% and 75% quarters show similar values for all files. Modified text was added in the beginning of the Results and Discussion section "From the *Opuntia* pool of transcripts, we analyzed three sizes of *query* files: 2,000, 6,000, and 10,000 transcripts in each input *query* file. Two experiments were executed. First 20 FASTA files (input

files for the workflow) for each query size were randomly created, see notebook “01 - Data Partitioning”. Each of these files were submitted independently as jobs to the batch systems on each cloud provider. For the second experiment, 120,000 transcripts were randomly selected and then partitioned in files with 2,000, 6,000, and 10,000 transcripts to analyze the relationship between query size, runtime and cost.”

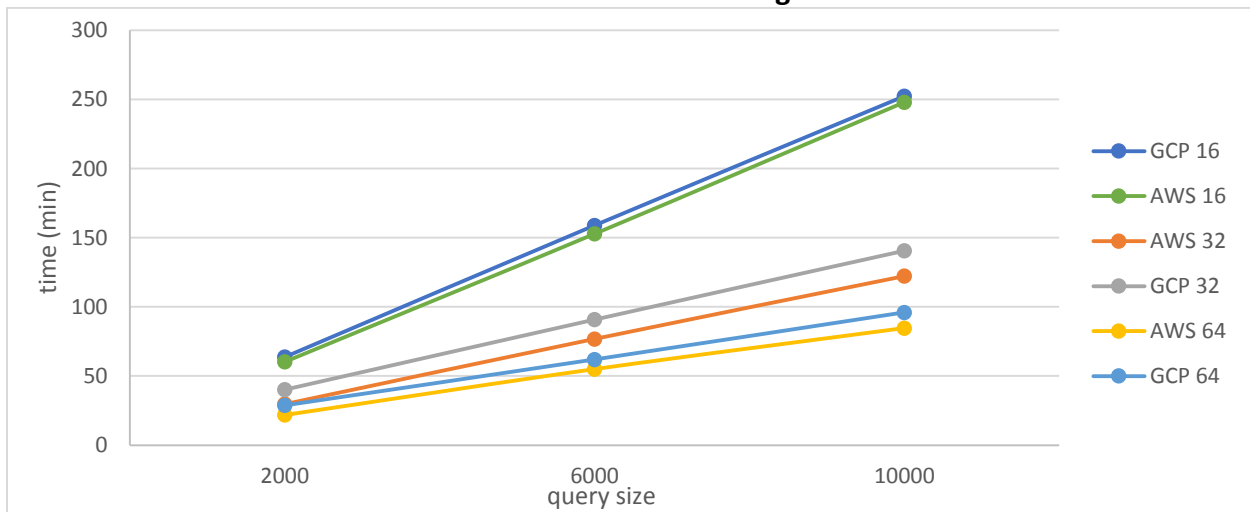
2. Please include a figure showing the distribution of transcript lengths for *Opuntia streptacantha*, and write that it is the prickly pear cactus. Presumably timing depends on the transcript lengths and on the number of BLAST hits.

Response: The referee’s comment about that the timing certainly depends on the transcript length and the number of BLAST hits. We added a cell to the 01 - Data Partitioning notebook that shows the transcriptome length distribution and its statistic metrics.

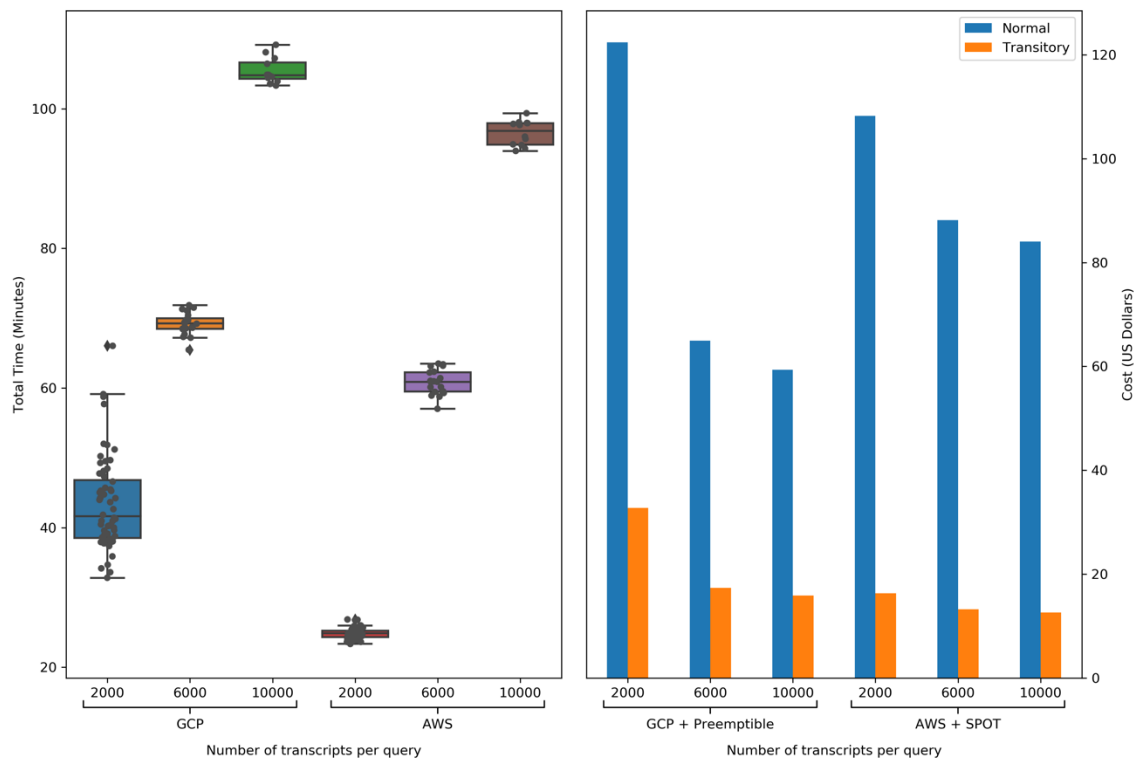
3. It is difficult to draw the conclusions from the way the data are plotted in Figure 4. The author concluded that "Reducing the number of transcripts per input file will reduce the total running time but will also increase the cost of the project as more instances will be in used." In addition to the raw data boxplot, it is better to show how the time and cost scale with query size, or with the number of instances. (The number of instances equals query size divided by total transcripts). Controlling the total transcripts and making instances as variable may be helpful in balancing the interpretation and data.

The plot below shows the relationship between the number of CPUs and time, and query size and time. (data collected using <https://automeris.io/WebPlotDigitizer/>). It also provides direct-viewing evidence to support the author's conclusion "AWS is more suitable for large data analysis groups to establish a set of queues and compute environments for multiple pipelines." Unlike the boxplot in the current manuscript, this figure also shows the differences in scaling between 16, 32 & 64 vCPU nodes.

Please add graphs showing query file sizes vs time (as below) and query file sizes vs cost. As well as cost vs time. These are the important take-home messages from the manuscript, but it is difficult to extract this information from the current figures.



Response: We agree with the referee’s comment that the relationship between query size, time and cost was not well described in the way that the results were presented in Figure 4. We add a new experiment where we processed 120,000 transcripts using the three query sizes. The new Figure 4 shows the relationship between the time, cost and query size for processing a fixed number of transcripts. Additional text describing this experiment was added to the Result and Discussion section: “Figure 4 shows the time and cost of processing 120,000 transcripts using second generation 64 vCPUs instances on each cloud provider. Reducing the number of transcripts per input file reduces the total run time but will also increase the cost of the analysis as more instances will be used. BLAST databases are transferred to more instances spending, on average, 10 minutes for each instance. For example, our experiment with the 10,000-query size processes all transcripts in about 105 minutes with a total cost of 59.37USD using 12 instances (GCP, N2, 64 CPU). Processing the same number of transcripts with a query size of 2,000 costs 122.36USD with all transcripts processed in 43 minutes using 60 instances (GCP, N2, 64 CPU).”



Optional (not necessary) suggestions:

The figures could be improved to give a clearer visualization of the data. For Figure 3a, will adding a secondary Y-axis regarding money and draw a line plot be better to show the relationship between total time and cost? For Figure 3b&c&d, drawing a component bar chart

could allow the readers to compare the job-dependent time among various configurations and demonstrates the proportionality at the same time.

Response: We thank the referee for this suggestion. In our opinion, as the data we are plotting here is not continuous data, drawing a line between points in Figure 3a could imply the idea of continuity. For Figures 3b, 3c, and 3d drawing bars could hide the variability of the runtime for each input file that is associated with instance and network performance. We prefer to keep the figure as it is now.

Reviewer #3: The authors provide a comparison of two cloud-based solutions for running BLAST-based transcriptomics analysis.

With cloud-based solutions becoming more popular in science, I think this comparison, along with the practical recommendations provided in this manuscript will be interesting to readers.

Response: We thank the referee for their comments and time dedicated to this manuscript.

Some suggestions for enhancements below:

1) The authors mention that there are numerous genomics companies in the space of cloud based biocomputations, citing reference 14 which discusses some of the legal responsibilities of groups doing such cloud based analyses. However, the authors do not connect this discussion back to the use of GCP or AWS, where users would need to obtain similar guarantees of data security which may not be possible to obtain. Thus a comparison of pricing against the private firms which provide similar services would be interesting to see if they provide specific guarantees that affect researchers with data that requires specific legal requirements.

Response: We thanks the referee for the comment. This study was executed under the NIH's STRIDES initiative using the currently available cloud provider partners: AWS and GCP. We do not consider that a comparison between public and private cloud providers and their legal responsibilities is within the scope of this study. However, we think that mentioning some of the most important private cloud providers is necessary. Accordingly, we rephrased the paragraph as:

In addition, private genomic cloud providers, for instance DNAnexus (www.dnanexus.com), DNASTAR (www.dnastar.com), Seven Bridges (www.sevenbridges.com) and SciDAP (scidap.com), also are in the market and offer cloud-based genomics frameworks. These commercial cloud providers make the execution of computational biology experiments easier offering command line and web-based interfaces designed for genomic data analysis.

2) I really like the inclusion of a best practices section with practical recommendations, and would love to see this section expanded:

a) In the first point the authors state: "We recommend CWL because the resulting product is portable and scalable, and it can be executed across a variety of computational environments as dissimilar as personal laptops or the cloud". However, these features are not exclusive to CWL, and solutions such as NextFlow and SnakeMake (and probably others) would also fit this description (and both of these also offer point 2 Conda and containerization). Please elaborate on your recommendation to include discussion of other workflow management systems, and explain in more detail why you would recommend CWL over these other solutions.

Response: We agree with the referee about the inclusion of more information about CWL and workflow managers. We added a new section to the manuscript named: **Common Workflow Language**.

b) In point 5, you recommend that users "Execute a small test in the cloud to find the best instance type for a workflow". Do you have any further practical recommendations about how users can best go about this? E.g. how does one define a "small test run" from a full datasets, and how can they predict how this will scale up to the full analysis and determine the most suitable machine types?

Response: We agree with the referee that defining a "small test run" may be difficult and it is intrinsically determined by the data and the workflow to be used. Our intention with this recommendation was to alert users that the cloud is a completely different environment than local workstations or on premise clusters. Users should test different cloud services and configurations before submitting a huge number of jobs. We edited that recommendation to:

5. *Cloud computing behaves differently than local workstations or on premise clusters. Users should define and execute small tests with their data and workflow before submitting large jobs. Testing different cloud services and configurations could help to reduce the runtime and cost for the whole analysis.*

3) It could be nice to expand the section about the Jupyter notebooks, and how these are being used, perhaps with some screenshots of results, and/or a small schematic showing that (if I understand correctly): the user interacts with the Jupyter notebook on their local machine, which in turn configures the cloud resources and starts the CWL workflow on the cloud, and then fetches the relevant results back and analyses them and displays results to the user. I think a schematic to this effect would be helpful for less technical readers and this in combination with some screenshots of the analysis results in Jupyter will increase the appeal of your work to research scientists.

Response: We agree with the referee about expanding the Jupyter notebook section. We added more description to it. The notebooks are available on Github for reading and browsing. Adding more figures to the manuscript would increase its size and complexity. Jupyter notebooks are

very popular and we think that less technical readers could easily find documentation about Jupyter notebooks without any problem.

4) In the conclusion the authors state “In our opinion, the choice of a cloud platform is not dependent on the workflow but, rather, on the specific details of the cloud provider”. However, I don’t believe the authors can make this statement having tested only a single workflow. So please rephrase the conclusion, or compare performance of different workflows covering a range of different characteristics (e.g. one that is memory-intensive, one that is CPU-intensive, and one that requires a lot of data transfer) and showing whether this conclusion holds, or whether some of the “specific details of the cloud provider” may make it more or less suitable for certain types of workflows.

Response: We agree with the referee that the phrase was general. We rephrased to this:

In our opinion, for BLAST based workflows, the choice of cloud platform is not dependent on the workflow but, rather, on the specific details and requirements of the cloud provider (e.g. NCBI maintains updated copies of the very large genetic sequence databases, such as nr, RefSeq and SRA, on both GCP and AWS). These choices include the accessibility for institutional use, the technical knowledge required for effective use of the platform services, and the availability of open-source frameworks such as application programming interfaces (APIs) to deploy the workflow.

5) AWS, Google, and Azure are probably the "big 3" providers that most readers will know about, and they might wonder why Azure was not included here and how it would compare. I understand that the authors cannot compare all providers, but it may be useful to at least mention Azure in the introduction where different providers are mentioned, and briefly explain if there were any specific reasons why you chose to compare AWS and GCP, and whether the same methodology could also be applied to Azure.

Response: We thank the referee for the comment. This study was executed under the NIH’s STRIDES initiative using the current available cloud provider partners: AWS and GCP. We don’t have access to Azure, therefore, we cannot extrapolate the conclusions of this study to Azure or any other cloud provider.

NIH Publishing Agreement & Manuscript Cover Sheet

By signing this Cover Sheet, the Author, on behalf of NIH, agrees to the provisions set out below, which **modify and supersede**, solely with respect to NIH, any conflicting provisions that are in the Publisher's standard copyright agreement (the "Publisher's Agreement"). If a Publisher's Agreement is attached, execution of this Cover Sheet constitutes an execution of the Publisher's Agreement, subject to the provisions and conditions of this Cover Sheet.

1. **Indemnification.** No Indemnification or "hold harmless" obligation is provided by either party.
2. **Governing Law.** This agreement will be governed by the law of the court in which a claim is brought.
3. **Copyright.** Author's contribution to the Work was done as part of the Author's official duties as a NIH employee and is a Work of the United States Government. Therefore, copyright may not be established in the United States. 17 U.S.C. § 105. If Publisher intends to disseminate the Work outside of the U.S., Publisher may secure copyright to the extent authorized under the domestic laws of the relevant country, subject to a paid-up, nonexclusive, irrevocable worldwide license to the United States in such copyrighted work to reproduce, prepare derivative works, distribute copies to the public and perform publicly and display publicly the work, and to permit others to do so.
4. **No Compensation.** No royalty income or other compensation may be accepted for work done as part of official duties. The author may accept for the agency a limited number of reprints or copies of the publication.
5. **NIH Representations.** NIH represents to the Publisher that the Author is the sole author of the Author's contribution to the Work and that NIH is the owner of the rights that are the subject of this agreement; that the Work is an original work and has not previously been published in any form anywhere in the world; that to the best of NIH's knowledge the Work is not a violation of any existing copyright, moral right, database right, or of any right of privacy or other intellectual property, personal, proprietary or statutory right; that where the Author is responsible for obtaining permissions or assisting the Publishers in obtaining permissions for the use of third party material, all relevant permissions and information have been secured; and that the Work contains nothing misleading, obscene, libelous or defamatory or otherwise unlawful. NIH agrees to reasonable instructions or requirements regarding submission procedures or author communications, and reasonable ethics or conflict of interest disclosure requirements unless they conflict with the provisions of this Cover Sheet.
6. **Disclaimer.** NIH and the Author expressly disclaim any obligation in Publisher's Agreement that is not consistent with the Author's official duties or the NIH mission, described at <http://www.nih.gov/about/>. NIH and the Author do not disclaim obligations to comply with a Publisher's conflict of interest policy so long as, and to the extent that, such policy is consistent with NIH's own conflict of interest policies.
7. **For Peer-Reviewed Papers to be Submitted to PubMed Central.** The Author is a US government employee who must comply with the NIH Public Access Policy, and the Author or NIH will deposit, or have deposited, in NIH's PubMed Central archive, an electronic version of the final, peer-reviewed manuscript upon acceptance for publication, to be made publicly available no later than 12 months after the official date of publication. The Author and NIH agree (notwithstanding Paragraph 3 above) to follow the manuscript deposition procedures (including the relevant embargo period, if any) of the publisher so long as they are consistent with the NIH Public Access Policy.
8. **Modifications.** PubMed Central may tag or modify the work consistent with its customary practices and with the meaning and integrity of the underlying work.

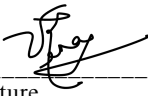
The NIH Deputy Director for Intramural Research, Michael Gottesman, M.D., approves this publishing agreement and maintains a single, signed copy of this text for all works published by NIH employees, and contractors and trainees who are working at the NIH. No additional signature from Dr. Gottesman is needed.

Author's name: Roberto Vera Alvarez, Leonardo Mariño-Ramírez, and David Landsman

Author's Institute or Center: NCBI/NLM/NIH Check if Publisher' Agreement is attached

Name of manuscript/work: Transcriptome annotation in the cloud: complexity, best practices and cost

Name of publication: GigaScience

Author's signature 

07/02/2020

Date