InfiniCloud: Leveraging the Global InfiniCortex Fabric and OpenStack Cloud for Borderless High Performance Computing of Genomic Data

Revision: 5 Oct 2015

RESPONSE TO REVIEWERS:

We are delighted that the reviewers found that our paper describing InfiniCloud interesting and we thank the reviewers for their helpful comments. We have made several revisions as suggested by the reviewers:

REVIEWER A

1. How does solution help understanding of mechanisms in genomics compared to other situations? Would it be economically efficient if implemented in other situations?

There are ongoing efforts internationally to sequence a large number of human genomes in order to catalog the genomic variation linked to diseases. The increased sample size will increase the statistical power of association tests that are used to discover hitherto unknown linkages between genomic variation and diseases. For example, Genomics England is currently targeting the sequencing and analysis of 100,000 genomes in order to understand the mechanisms behind the genetic basis of diseases. The current bottleneck with the increasing adoption of genomic sequencing is in the data transfer and processing of the genomic data, which can be addressed with the ability of computing resources to be scaled out beyond any one site when resources are not readily on hand.

For the proof-of-concept test of long range InfiniBand transport, we have revised the text to emphasize its utility in the scaling up of compute resources beyond a single site by utilizing spare fibre capacity on existing networks (sections 2.1.3 and 3.1) and linking compute resources at remote sites.

- The transfer of data from Australia to Singapore was a proof-of-concept experiment to demonstrate that the native InfiniBand transport would allow for one site to scale beyond local capacity and utilize resources at remote sites. In this case, genomic data was transferred from Australia to Singapore in one step of a multi-stage bioinformatic pipeline to take advantage of a large memory system that was not readily available in Australia at that point of time. The rest of the pipeline was executed in Australia using the computed data.

- The native InfiniBand transport can be implemented on spare fibre capacity on existing networks, which obviates the need for laying new cables. We have included a clarifying statement in section 2.1.3.

2. Details about data

We have included details about the genomic data in the reference (Section 3.2; reference 8).

3. Discussion about transfer of human genome datasets over 100M network

We have included a short discussion about the time needed for transfer of genomic dataset in section 4 (paragraph 3) using the output from a single Illumina HiSeq X machine as an example. Briefly, the output from 32 genomes could be transferred over a dedicated TCP/IP 100M network in a week assuming maximal bandwidth with 100% efficiency. In practice, it is difficult to achieve anything close to line rate and this would preclude the scaling up of genomic data transfers from additional sequencing machines. In contrast, the native InfiniBand tests under real world conditions show at least an order of magnitude improvement in transfer rates.

4. Distracting shell output

We have removed certain sections of the shell output and formatted them for clarity.

5. Description of VMs and their use in genomics

We have included a description of virtual machines and expanded on the explanation of its usefulness in genomic workflows (Section 1 paragraph 3).

6. Focus on how proposed hardware solution would help the study of genomics data

We thank the reviewer for the suggestion, and we have removed most of the shell output to focus on how InfiniCloud was used in the study of genomic data.

7. Editing of title

We have omitted the phrase '..and beyond' from the title, but we feel that inclusion of the phrase 'genomic data' would be appropriate to set the context for how InfiniCloud was used in this paper.

REVIEWER C

In the presented article the authors say about the transfer huge files from NCI to A*STAR and vise versa. For example, the analysis of genomic data require to transfer about 233 GB date from Singapore to Australia (page 13). In this case, the main role is played by the bandwidth, rather than delay. However, no comparison their results and similar results of TCP/IP data transfer using FTP and GridFTP is presented.

The comparison between long-range InfiniBand (dsync+) and TCP (rsync) has been performed in reference [10]. This has been added in section 3.2 paragraph 2.

1. The list of keywords is absent.

The list of keywords has been added to the abstract.

2. Fig. 4 is identical to Fig.2.

We have removed the duplicate figure.

3. The references to Figures 2, 3 and 6 in the text are absent.

We have included the references to the figures in the revised text.

4. The reference to Table 4 in the text is absent.

We have included the reference in the revised text.

5. Fig. 5 is too small to read. In black and white printing mode it does not understand what is what.

The above defects have been corrected in text.

6. All listings should be decorated by box and supplied the captions.

The remaining listings are now decorated by boxes and are supplied with captions

7. Some listings should be removed because do not contain useful information or extra information. For example, the output of ifconfig command (page 6), the listing demonstrated bandwidth capability (page 6), the listing on page 7, the listing is on pages 8-10, the output of ls command (page 13) and the listing is on the page 14.

8. Remove Fig.6.

We have removed the figure

8. Page. 9. No references to the SGE job scheduler, Ganglia monitoring tools, and the IPython notebook shell which mentioned in Fig. 7 too.

We have included references to the tools used in the ElastiCluster setup.

9. Both tab. 6 and all pieces of main text should be carried out from the

listing.

The above defects have been corrected in text.

REVIEWER D

1. Perhaps, the size of the annotation (approx. 400 words) could be somewhat reduced, leaving a brief description of the results. It would also bring the article into line with the requirements to the authors of the journal (Abstract's optimal size is 100 to 150 words)

We have omitted some of the figures (e.g. shell output) in the paper to reduce the size of the results section. The abstract has also been shortened.

2. There is a statement that: A major part of the setup effort was for OpenStack to play nicely with the InfiniBand interfaces. For this, a few OpenStack out-of-tree patches were necessary such as forcing use of one network PKEYS. It would be interesting to get a more detailed description of the OpenStack "tuning process" to work with InfiniBand.

We have provided additional details of the implementation in Section 2.2 “InfiniCloud Installation and Configuration”, paragraph 3.

3. There is a question, whether it was possible to obtain any test results on systems without the SR-IOV technology? If so, they would allow to show advantages of this approach compared with older methods of networking organization.

We have addressed this in Section 2.2 “InfiniCloud Installation and Configuration”, paragraph 2. As the implementation of native Infiniband transport in virtual instances can only be done with SR-IOV technology, we are not able to perform comparable tests.