Improving the scalability of movement monitoring workflows: An architecture for the integration of the Hadoop File System into e-Science Central

Hugo Hiden[1], Simon Woodman[1], Paul Watson[1], Michael Catt[2], Michael Trenell[3]
[1]School of Computing Science
[2]Institute for Ageing and Health, The Medical School
[3]Movelab, The Medical School
Newcastle University
Newcastle upon Tyne
{hugo.hiden|simon.woodman|paul.watson|michael.catt|michael.trenell}@newcastle .ac.uk

Shaoyan Zhang
Unilever Discover
Colworth Science Park
Bedford
MK44 1LQ

Understanding patient activity levels is important to assessing key lifestyle variables linked to obesity, diabetes and cardiovascular disease. The MOVEeCloud project [1] makes use of wrist worn accelerometers to measure movement data over three axes at approximately 80Hz. Once collected, the analysis procedure for this data involves categorizing the acceleration signals into one of several categories: Sedentary, Light Activity, Walking and Running [2]. As each data file collected is approximately 2GB in size and comprises some 100 million rows of data, this filtering process is time consuming (taking approximately one hour to process a single sample). However, because the categorization algorithm breaks the data into short chunks (1024 samples, representing 12.8 seconds of data), these files can be broken into chunks and processed separately. Once processed, results are presented in the form shown below in Figures 1 & 2.

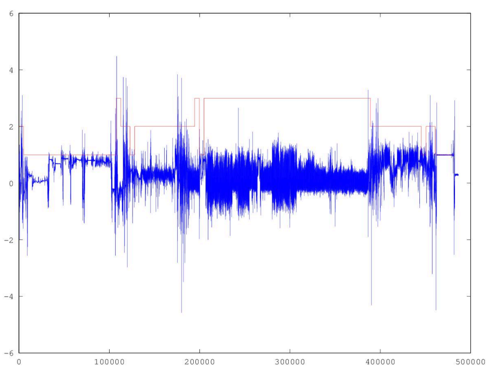

Figure 1: Categorized accelerometer data

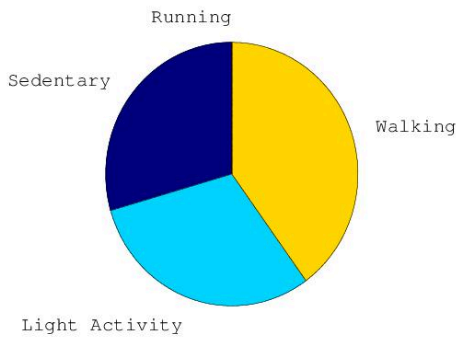

Figure 2: Categorization results

In order to process the data, researchers have developed code in Matlab, which reads data files and generates a summary chart and counts of the various activity levels. This code was converted into Octave [3] and deployed within the e-Science Central (e-SC) workflow engine [4]. Each set of data is first broken into chunks using a workflow which splits the data into blocks representing one day of data collection (Figure 3). Then, for each of these chunks a categorization workflow is executed (Figure 4). These categorization workflows can execute in parallel with the top level workflow blocking until all of the chunks have been processed.

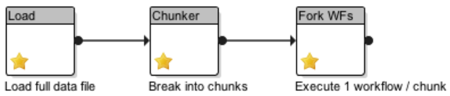

Figure 3: Top level "chunking" workflow

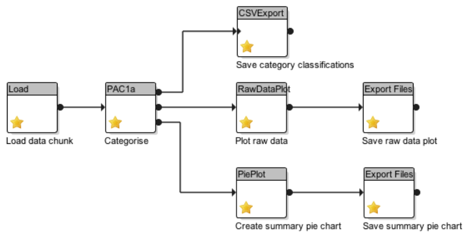

Figure 4: Chunk Processing workflow

Whilst this approach worked for other projects, the design of e-SC is such that data supplied to workflows is supplied by a single source – the main server. Although excellent scalability results have been reported using this technique [5], it was found that as the number of workflow engines was increased beyond 100 this single point of access to data introduced a severe performance bottleneck . This bottleneck is particularly significant as the size of data files transferred in this project is several orders of magnitude greater than those encountered during the drug discovery study.

There are several approaches that could have been adopted to address this bottleneck - One would have been to allow workflow blocks to store and retrieve data directly from Azure Blob Store or Amazon S3. This has been tried and shown to work well [5], however this approach is highly cloud specific with data storage then residing outside of the e-SC system. This means that the provenance capture and security facilities provided by e-SC are not available. Also, using cloud storage directly does nothing to address the underlying scalability bottleneck of the e-SC workflow system. The technique adopted in this study, therefore was to use the individual workflow enactment nodes as a data storage infrastructure for e-SC and deploy the Apache Hadoop File System (HDFS)[6] to manage this storage.

HDFS is the underlying storage system for the Apache Hadoop [7] framework which allows distributed processing of large data sets across large numbers of commodity servers. The HDFS component of Hadoop provides a robust storage cluster which can be used to persist large data sets and provide scalable access to thousands of clients. It is the scalability and replication features of HDFS that make it an attractive option to use as the underlying storage system for this movement data analysis application.

In order to integrate HDFS, we created a new storage driver for e-SC (drivers already exist for Azure, S3 and plain filesystems) and provided the API client used by the workflow engines with direct access to this storage. In addition, each workflow enactment node was configured as an HDFS data node. To support the HDFS filesystem, an additional Virtual Machine was instantiated to act as the master Name Node (Figure 6). The entire infrastructure was deployed as Linux Virtual Machines within the Microsoft Azure platform [8].

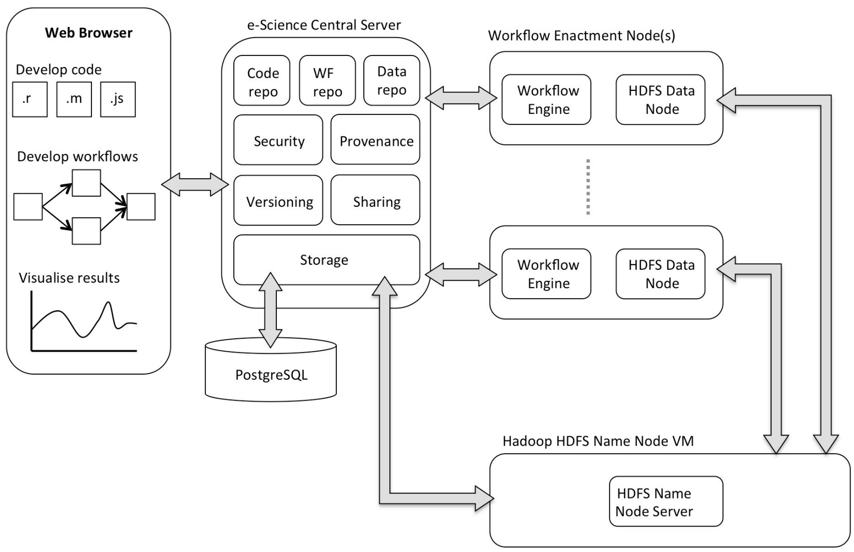

Figure 6: System Architecture

Using this system allows workflow engines to access data directly from HDFS, but it does not use the location awareness built into HDFS to execute workflows on nodes that physically contain the data. This is due to the current design on the e-SC workflow enactment subsystem that routes workflow execution requests via a JMS queue to the first available enactment node. It is anticipated that finer grained control over this workflow request routing would lead to increased performance gains. As this work is still underway a full set of performance data has not yet been collected, but data will be presented that demonstrates a significant performance improvement over the basic e-SC data storage architecture.

References

1. MOVEeCloud homepage. http://movelab.org/research/moveecloud/
2. Zhang S, Rowlands AV, Murray P, Hurst TL. (2012): "*Physical activity classification using the GENEA wrist-worn accelerometer*". Med Sci Sports Exerc. 2012 Apr;44(4):742-8.
3. The GNU Octave project: http://www.gnu.org/software/octave/
4. Hiden, H.G., Woodman, S.J., Watson, P., Cala, J. (2012): "*Developing Cloud Applications using the e-Science Central Platform*". To appear in Philosophical Transactions of the Royal Society A.
5. Cala, J., Hiden, H.G., Watson, P., Woodman, S.J. (2012): "*Cloud computing for fast prediction of chemical activity*". 2nd International Workshop on Cloud Computing and Scientific Applications(CCSA). May 13-16 2012, Ottawa, Canada.
6. Apache Hadoop Distributed File System. http://hadoop.apache.org/hdfs/
7. Apache Hadoop project. http://hadoop.apache.org/
8. Microsoft Windows Azure homepage. http://www.windowsazure.com