

# Cloud computing in the age of data-intensive science

G. Bruce Berriman, Ewa Deelman, and Gideon Juve


*Internet-based computing may prove a powerful tool for generating new science products, and emerging software technologies hold great promise for capturing records of their processing history.*

The late Jim Gray<sup>1</sup> described how science is becoming ‘data-centric,’ in the sense that vast quantities of public data are used to create new products. What is the cheapest and most efficient way of producing them? Should the data be uploaded to the software, or should software be built near the data? How can we curate these products and capture their processing histories?

We have investigated two aspects of these broad themes.<sup>2</sup> To examine the cost and efficiency of generating new products, we created new image sets from data uploaded to the Amazon Elastic Compute 2 (EC2) cloud and compared its performance with that of the Abe cluster at the National Center for Supercomputing Applications. We then explored the use of execution logs to capture the processing history of these products.

We used the Montage image-mosaic engine<sup>3</sup> for image processing, which creates a composite from multiple images. We computed all mosaics in three stages. First, the input images were reprojected to distribute the energy in the input pixel pattern on the sky to the output pixel pattern. Next, the sky background radiation in each reprojected image was rectified to a common level across each image. Finally, the reprojected, rectified images were co-added to produce the final mosaic. The output of one process becomes the input to the next. Thus, Montage is a data-driven workflow or pipeline application. Because it spends 95% of its time on input/output (I/O) operations, it is described as I/O-bound.

We generated eight-square-degree image mosaics of the Messier 17 star-forming region based on 4GB of Two Micron All Sky Survey (2MASS) images. The workflow contained over 10,000 tasks and produced 8GB of output data. Our goal was to compare the performance of Amazon EC2 and Abe. Because the former uses commodity hardware while the latter operates on high-speed networks, we generated all mosaics on single nodes,



**Figure 1.** Performance of Montage, Broadband, and Epigenome on various platforms. The legend identifies the processors (see Table 1). Processors designated ‘m1’ and ‘c1’ are on the Amazon EC2 cloud, while those designated ‘abe’ are installed in the Abe cluster.

employing Intel chips running Red Hat linux, to allow a side-by-side comparison.

We managed all jobs with the Pegasus workflow-management system,<sup>4</sup> which transforms high-level descriptions of workflows into specific sequences of operations and identifies the computing resources required for execution. It then transfers control to standard grid-based management tools. We compared the Montage performance to that of Broadband, a memory-bound seismology simulation application, and Epigenome, a processing-bound biochemistry application. We also compared the costs of owning and maintaining hardware locally with those of Amazon EC2, which charges for provisioning and running processors and for uploading, storing, and downloading data, and returns itemized invoices for each job. Figures 1 and 2 show the performance and cost, respectively, of running Montage, Broadband, and Epigenome on the processors installed on Amazon EC2 and Abe (see legends and Table 1).

On Amazon EC2, the best performance for Montage was obtained using the resource with the most memory (m1-xlarge)—see Figure 1—likely because of superior file-system cache performance. However, the parallel file system on Abe offers the best

**Table 1.** Characteristics of the Amazon EC2 (m1 and c1) and Abe platforms (abe) used to derive the results in Figures 1 and 2. Arch: Architecture. CPU: Central processing unit. InfiniBand is an input/output technology that connects high-performance processors with high-speed storage devices via bidirectional serial links.

Type	Arch	CPU	Cores	Memory	Network	Storage	Price
m1.small	32bit	2.0–2.6GHz Opteron	1	1.7GB	1Gb/s Ethernet	Local	\$0.10/hr
m1.large	64bit	2.0–2.6GHz Opteron	2	7.5GB	1Gb/s Ethernet	Local	\$0.40/hr
m1.xlarge	64bit	2.0–2.6GHz Opteron	4	15GB	1Gb/s Ethernet	Local	\$0.80/hr
c1.medium	32bit	2.33–2.66GHz Xeon	2	1.7GB	1Gb/s Ethernet	Local	\$0.20/hr
c1.xlarge	64bit	2.0–2.66GHz Xeon	8	7.5GB	1Gb/s Ethernet	Local	\$0.80/hr
abe.local	64bit	2.33GHz Xeon	8	8GB	10Gb/s InfiniBand	Local	...
abe.lustre	32bit	2.0–2.6GHz Opteron	8	18GB	10Gb/s InfiniBand	Parallel	...

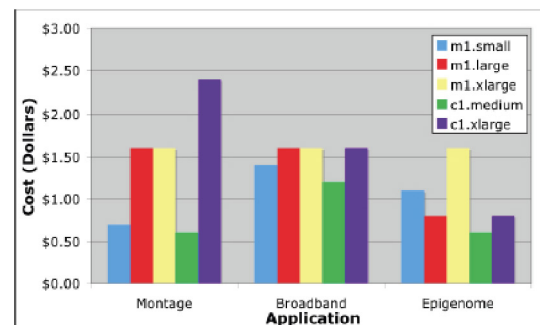
performance for I/O-bound applications such as Montage. For Broadband (memory bound) and Epigenome (processor bound), this advantage disappears.

We consider the processing and data costs for Amazon EC2 only. There is no need to choose the most powerful processor to get the best-value performance for I/O-bound applications. The c1.medium processor, a 32bit machine with only 1.7GB of memory, offers the best value without appreciable loss of performance. The most powerful processors come into their own for Broadband and Epigenome, where they provide the best performance at reasonable cost.

Because Montage is data-intensive, it incurs higher costs in uploading, storing, and transferring data than Broadband and Epigenome. On the most cost-effective processor, Montage data costs are \$1.75 per job, which is higher than the processing cost of \$0.50. For the other applications, the corresponding data and processing costs are \$0.40 and \$0.60 per job, respectively.

To address our third question, we augmented Pegasus with an existing open-source provenance methodology, the Provenance Aware Service Oriented Architecture (PASOA), which is already used in fields such as aerospace engineering.<sup>5</sup> When Montage runs, Pegasus creates a provenance record in eXtensible Markup Language (XML) that describes the processing history. This record can be stored in a database to create a permanent, searchable store. Although this process is experimental, the first results are encouraging.

While the Amazon EC2 cloud is at a performance disadvantage relative to a high-speed cluster for I/O-intensive applications, it is a powerful tool for science-workflow applications. Under current cost models, Amazon EC2 data costs are relatively high, but the total costs of running a job are comparable to those incurred on a local server. Amazon EC2 is more cost-effective for memory- and processor-bound than for I/O-bound applications.



**Figure 2.** Processing costs of running Montage, Broadband, and Epigenome on various platforms on Amazon EC2. The legend identifies the processors (see Table 1).

We plan to extend our work by studying astronomical applications, such as processing-bound periodogram calculations that compute the power in the periodicities present in time-varying data sets. We performed the current study on single nodes, but we intend to examine performance optimization by employing multiple nodes simultaneously and analyzing ways to distribute data between nodes.

G. Bruce Berriman acknowledges support from the Jet Propulsion Laboratory at the California Institute of Technology under contract with NASA. Ewa Deelman and Gideon Juve acknowledge support from the National Science Foundation under the SciFlow (CCF-0725332) and Pegasus (OCI-0722019) grants. This research made use of Montage, funded by NASA's Earth Science Technology Office, Computation Technologies Project, under cooperative agreement NCC5-626 between NASA and the California Institute of Technology.

Continued on next page

## Author Information

---

### G. Bruce Berriman

Infrared Processing and Analysis Center  
California Institute of Technology  
Pasadena, CA

Bruce Berriman obtained his PhD from the California Institute of Technology. His research interests include using archived data to search for new brown dwarfs and application of new technologies to scientific computing.

### Ewa Deelman and Gideon Juve

Information Sciences Institute  
University of Southern California  
Marina del Rey, CA

Ewa Deelman is a research associate professor in the Computer Science Department. Her research interests include design and exploration of collaborative, distributed scientific environments.

Gideon Juve is a PhD student in computer science. His research interests include distributed and high-performance computing, scientific workflows, and computational science.

## References

1. T. Hey, S. Tansley, and K. Tolle, *Jim Gray on eScience: a transformed scientific method*, in T. Hey, S. Tansley, and K. Tolle (eds.), **The Fourth Paradigm: Data Intensive Scientific Discovery**, Microsoft Research, Redmond, WA, 2009.
2. G. Juve, E. Deelman, K. Vahi, G. Mehta, B. Berriman, B. P. Berman, and P. Maechling, *Scientific workflow applications on Amazon EC2*, in **IEEE Cloud Comput. Wkshp**, 2009.
3. <http://montage.ipac.caltech.edu> Montage web page.
4. <http://pegasus.isi.edu> Pegasus workflow-management-system web page.
5. G. B. Berriman and E. Deelman, *The role of provenance management in accelerating the rate of astronomical research*, **Proc. Science**. In press. First presented at 'Accelerating the Rate of Astronomical Discovery,' 2009, Int'l Astron. Union Gen. Assembly, Spec. Sess. 5, Brazil.