An Overview of the Open Science Data Cloud

Robert L. Grossman*
University of Illinois at Chicago

Yunhong Gu
University of Illinois at Chicago

Joe Mambretti
Northwestern University

Michal Sabala
University of Illinois at Chicago

Alex Szalay
Johns Hopkins University

Kevin White
University of Chicago

ABSTRACT

The Open Science Data Cloud is a distributed cloud based infrastructure for managing, analyzing, archiving and sharing scientific datasets. We introduce the Open Science Data Cloud, give an overview of its architecture, provide an update on its current status, and briefly describe some research areas of relevance.

Categories and Subject Descriptors

C.1.4 [**Parallel Architectures**]: Distributed architectures; C.2.4 [**Distributed Systems**]: Distributed applications; H.2.8 [**Database Applications**]: Scientific databases; H.3.4 [**Systems and Software**]: Distributed systems

General Terms

Design, Experimentation

Keywords

cloud computing, data clouds, management and analysis of scientific data

1. INTRODUCTION

Many scientists today find it a challenge to manage and analyze their data. There is a growing *gap* between the ability of modern scientific instruments to produce data and the ability of scientists and most research groups to manage, analyze, and share the data that is produced. Indeed, as the size of datasets grows from MB to GB and TB, many scientists are finding it difficult just to transport their data from the instrument where it is produced to the facility where it is analyzed, and from the facility where it is analyzed to their collaborators around the world.

*The authors are listed in alphabetical order. Robert Grossman is also an Associate Senior Fellow at the Institute for Genomics and Systems Biology at the University of Chicago.

This is an exciting time of change for infrastructure that supports data intensive science [6]. See Figure 2. First, there is an increasing demand for more sophisticated data analyses over datasets of ever increasing sizes. Second, the limitations of the current databases and grids for data intensive science are becoming painfully clear and new computational architectures are being developed [6], large data clouds [12, 13, 21, 20], pioneered by Google, GrayWulf clusters pioneered by Szalay and Microsoft [28], and new database architectures, such as SciDB being developed by Stonebraker et. al. [30] and so-called No-SQL databases, such as Hive [32]. Third, there are new computing patterns being introduced, including those for on-demand computing [2, 14], distributed computing [25], and integrating computing with data mining [5, 18].

As Figure 1 illustrates, it is helpful to think of data as coming in three sizes: small datasets, medium to large datasets, and very large datasets. By a small dataset, we mean a dataset small enough that it is easy for an individual scientist to manage with a single workstation. By a very large dataset, we mean a dataset large enough that a specialized data management infrastructure is required. A good example is the LHC Computing Grid developed for the Large Hadron Collider [23]. The challenge is that many medium to large datasets are becoming more common and are to difficult for an individual researcher or research group to manage themselves.

In this article, we introduce the Open Science Data Cloud or OSDC. The OSDC was started in 2010 and is a persistent, distributed storage and computing resource designed to manage, analyze, share, and archive scientific data. It is especially useful for medium to large datasets.

The OSDC is a hosted platform managed by a single entity (the Open Cloud Consortium), not a federated or virtual organization. Once data is accepted by the OSDC, the goal of the OSDC is to manage and provide access to the data from that time forward.

As we describe below in Section 3.2, our assumption is that if the OSDC invests a constant dollar amount each year to buy new equipment, then we can maintain all the old data and, indeed, add new data, because the capacity of disk doubles approximately every 18 months (or less). For example, if: 1) disk capacity improves by 50% each year, 2) 10 new racks are added each year i that provide D_i PB of storage, and 3)

10 old racks are retired with capacity D_{i-3} PB, then every three years the *additional capacity* $D_i - D_{i-3}$ provided by the 10 new racks is larger than the size of the entire OSDC collection in year $i - 3$.

The Open Cloud Consortium is managed by a not-for-profit, which has the goal of raising funding or acquiring donated equipment so that this constant acquisition of new equipment can be maintained and the integrity and access to data can be provided over long periods, not just within the scope of a project, which is often just 2-5 years.

Currently, storage for medium to large datasets for the OSDC is provided by the Hadoop [21] or Sector [20] large data clouds. The OSDC also supports elastic, on demand virtual machines, similar to Amazon's EC2 service [3]. In addition, the applications we describe below integrate traditional relational databases with these cloud services.

The different data centers in the OSDC are all connected by 10 Gbps networks. In addition the OSDC has a 10 Gbps link to the StarLight facility in Chicago [29], which in turn connects to the National Lambda Rail, Internet2, and a wide variety of research networks. With these connections, medium to large datasets can easily be ingested and shared. In addition, these connections are used so that medium and large datasets can be replicated over multiple geographically distributed data centers.

The OSDC is different from existing cloud resources today:

- It is designed to provide long term persistent storage for scientific data, even large scientific datasets.
- It can utilize high performance research networks, so that medium to very large datasets can be easily ingested, accessed, and shared over wide areas.
- It is a balanced architecture that uses data locality to support the efficient execution of queries and analysis over the data managed by the cloud.

In this article, we introduce the Open Science Data Cloud, give an overview of its architecture, provide an update on its current status, and briefly describe some research areas of relevance.

2. RESEARCH CHALLENGES

There are several challenges when working with large datasets: managing the data, analyzing the data, archiving the data, and transporting the data. As we will describe in the following section, the approach of the OSDC is to use a specialized large data cloud to handle these challenges. In this section, we briefly introduce these challenges and provide some background on clouds.

Transporting large datasets. Although most universities today have access to wide area, high performance networks, most scientists do not. There are two main problems: first, the laboratories of most scientists are not connected to high performance networks, although their universities are. Second, TCP as it is commonly deployed, does not use the bandwidth available on high performance networks effectively when large datasets are transported over wide areas [10].

Architectures that balance data management and data analysis. Although databases are highly optimized for managing and accessing indexed relational data, there are not optimized for numerically intensive operations on subsets of the data that comprise a significant portion of the data under management. That is, as the numeric computations require touching more and more of the data, it is more efficient to perform the computation outside of the database. In contrast, systems such as Hadoop are designed to perform numerically intensive computations on large datasets, especially computations that require scanning most of the dataset. Hadoop accomplishes this by exploiting data locality to colocate the computation over the data whenever possible. It is an important open research problem to develop systems that can exploit data locality for a wider set of data intensive applications than currently handled by Hadoop.

Archiving digital data. It is a difficult challenge to keep digital data for long periods of time. Problems include: moving to new formats, migrating to new media, refreshing current media, and losing the ability to meaningfully interpret the data that is archived. For more information about some of the problems, see [7].

Cloud Computing. Cloud computing does not have a standard definition yet, but a commonly used one has been developed by NIST [24] and defines cloud computing as "a model for enabling convenient, on-demand network access to a shared pool of configurable computing resources that can be rapidly provisioned and released with minimal management effort or service provider interaction." For the purposes here, we are interested in cloud computing platforms that not only provide on demand computing instances but also support data intensive computing, even over very large datasets. From this point of view, we are interested in two different but related architectures: the first type of clouds provide on-demand infrastructure, such as on-demand computing instances [4]; the second type of clouds provide scalable, on-demand services, such as storage or compute services. Both use similar machines, but the first is designed to scale out by providing additional computing instances, whereas the second is designed to support data-intensive or compute-intensive applications via scaling capacity.

Amazon's EC2 computing instances [1] are an example of the first type of cloud. The Eucalyptus system [33] is an open source implementation that provides on-demand computing instances and shares the same APIs as Amazon's EC2 cloud. Google's Google File System, BigTable and MapReduce infrastructure [12], [16] is an example of the second type of cloud. The Hadoop file system (HDFS) and Hadoop's MapReduce implementation [8] is an open source implementation of the Google cloud. A similar system is Sector/Sphere, which consists of the Sector distributed file system and Sphere computing service, which supports User Defined Functions (UDFs) over the data managed by Sector [20]. The Google stack, Hadoop and Sector/Sphere all

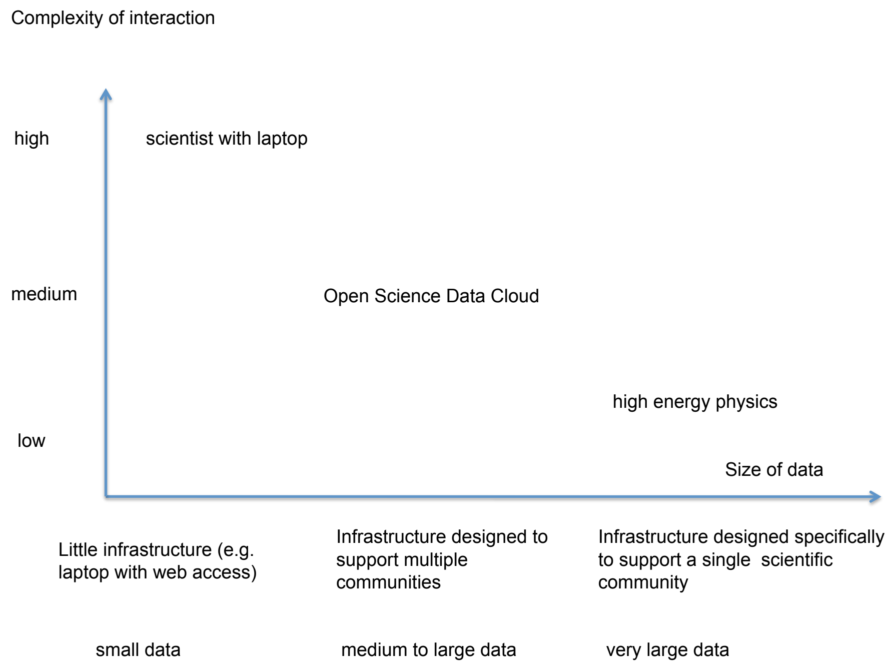

Figure 1: The Open Science Data Cloud is designed to be an infrastructure that can not only be used by multiple scientific communities, but also allows federated queries, correlations, and other analyses that span data sets from multiple scientific communities.

support data intensive computing by exploiting locality so that, to the extent practical, computation is done in-place over data managed by a scalable distributed file system. For a good comparison between cloud computing and grid computing, see [15].

3. ARCHITECTURE

3.1 Hardware Architecture

Based upon experimental studies that we have done over the past year [17], the hardware for the OSDC consists of commodity computers, networked with commodity 1GE connections that are balanced so that there is one disk for each core. To say it another way, we are currently balancing the system by allocating *one spindle per core*. A system that was more compute intensive would generally have more cores and fewer spindles. A system that was used mainly for storage would generally have fewer cores and more spindles. Notice that this architecture that uses one spindle per core is different than the standard architecture for a database, which usually consists of an upper end workstation connected to a RAID or storage area network.

The current generation of racks that we are using for the OSDC cost less than $90,000 each and consist of 32 nodes containing 124 cores, 496 GB of RAM and 124 1-TB disks. The rack consumes about 10.3 KW of power (excluding the power required for cooling). We have been calling these Raywulf Racks.

With 3x replication, there is about 40 TB of usable storage available, which means that the cost to provide balanced long term storage and compute power is about $2,000 per TB. So, for example, a single rack could be used as a basis for a private cloud that can manage and analyze approximately 40 TB of data.

Each rack is a standard 42U computer rack and consists of a head node, 31 compute/storage nodes, a top of the rack switch with two 10 Gbps uplinks that connect eventually to StarLight, and a separate switch for managing the servers that connects to the commodity Internet. We installed GNU/Debian Linux 5.0 as the operating system. In the racks that we are currently using, each compute/storage node is an Intel Xeon 5410 Quad Core CPU with 16GB of RAM, a SATA RAID controller, four (4) SATA 1TB hard drives in RAID-0 configuration, and a 1 Gbps NIC. The top of the rack switch is a Force10 S50N switch. Each rack also contains a 3Com Baseline 2250 switch so that we can manage the servers remotely over the commodity Internet using the Intelligent Platform Management Interface (IPMI) interface in each server.

We expect that the configuration of the racks will change from year to year and that over time we will migrate to a container-based instead of a rack-based infrastructure. For example, the next racks we buy will use 2 TB disks, and, in another year or two, we expect to use to 10 Gbps NIC.

3.2 Migration Architecture

In the initial phase of the OSDC, we plan on spending approximately an equal amount of capital each year to acquire new racks and to retire racks after approximately three years. With a plan like this, and assuming we start with 10 racks the growth of the OSDC can be seen in Table 1.

We anticipate that over time we may replace disk with solid state media or whatever emerges as a replacement for disk. We will use the ability of Hadoop and Sector (and whatever emerges to replace them) to replicate automatically the data so that there are always replicas available within the same data center and in at least one other data center.

In more detail, each year we will invest in new equipment. Let C_i be the amount of data (in PB) managed by the OSDC in year i and D_i the amount of new disk (in PB) that we add each year. Then we assume that we will retire D_{i-3} disk each year. Our assumption is that

$$D_i - D_{i-3} > C_{i-3}. \quad (1)$$

Since Hadoop and Sector both replicate data *automatically* when a rack is removed and when a new rack is added, the data refresh problem is managed through the strategy of adding sufficient new disk (or other appropriate media recognized by the software) so that Equation 1 holds.

Part of the Governance plan for the OSDC is to select new data to add to the facility and to make sure that the amount of data added each year is no greater than the net new capacity available each year.

3.3 Software Architecture

The OSDC software stack includes: the Hadoop [21] Distributed File System and Hadoop's implementation of MapReduce; and the Sector system [20] that provides a wide area storage system and an implementation of both MapReduce and cloud-wide User Defined Functions.

UDT [19] is used by Sector for communicating between the nodes, both within and across the data centers. As is well known [10], it can be challenging to transport large datasets using TCP using high performance networks over wide areas. UDT is a UDP-based protocol that implements reliability, can effectively utilize the bandwidth available on wide area high performance networks, is fair to multiple large data flows, and is friendly to standard TCP traffic [19].

Both Hadoop and Sector exploit data locality to improve the performance of data intensive computations. There are some interesting differences between how Hadoop and Sector support data locality. See Table 2. An important area of research for the OSDC is to improve our understanding of how to build scalable systems that support data locality for a variety of different types of data intensive applications.

Unlike the majority of current cloud deployments, the OSDC was created specifically to address the requirements of high volume data flows, such as the large streams used by data intensive science. Most clouds today were designed for millions of small data flows, each processing a small amount of information, supported by commodity Internet technologies. Much science research requires the processing of extremely large amounts of data that cannot be supported by generally implemented commodity networking technologies. Large scale data intensive science requires high performance support for both extremely high capacity individual data streams and for flexible interactive provisioning of such streams. Consequently, a key objective of the OSDC design was to create capabilities by using libraries like UDT that

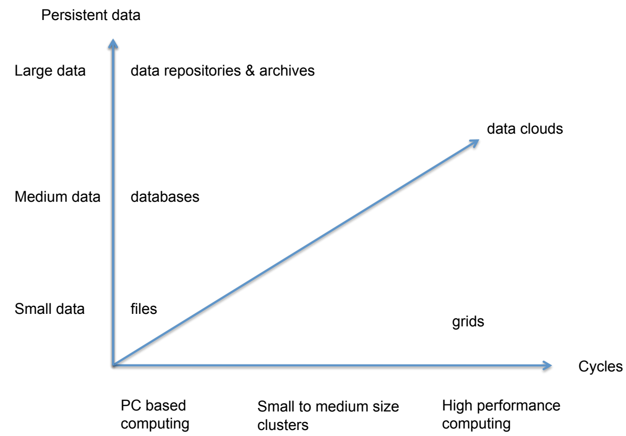

Figure 2: Data clouds are designed to provide long term, persistent storage for data and to provide co-located storage so that computation can be done over the data in a *balanced* fashion that scales from small data to large data. In contrast, grids and high performance computers as usually designed do not co-locate computing and storage. Instead, data is shipped to high end computing resources, placed in queue, and data is returned after the computation completes.

not only provide support for large data flows but also can be used with systems that dynamically provision networking resources.

The OSDC also include the Eucalyptus system [33], which provides elastic on-demand computing instances. The virtual machine images that we use are compatible with the Amazon public cloud, and so, if the data is small enough to be transported to an Amazon storage service, then it is easy to interoperate the OSDC private cloud with the Amazon public cloud.

The OSDC also includes three databases (MySQL, PostgreSQL and SQL Server) that are integrated with the cloud, as well as a variety of open source software for scientific computing, including R and Octave. In addition, the OSDC includes a GrayWulf style federated database as described in [31].

4. OSDC APPLICATIONS

4.1 Cistrack

Cistrack is an OSDC application to manage and share cis-regulatory data for the fly and worm produced by the NIH modENCODE Project [9]. Much of the data is created by next-generation sequencing machines that produce raw datasets that are a TB or larger in size and intermediate datasets that are 10 GB and larger in size. Cistrack already has more than 2 TB of data in a MySQL database and over 20 TB of data in a Sector archive. We make this data available for bulk download via the OSDC.

In addition, the OSDC provides the ability to analyze the data with user-supplied pipelines that are executed using elastic on-demand virtual machines managed by Eucalyptus.

Year	New Racks	Total Racks	New Cap.	Total Cap.	Net New Cap.
1	10	10	1.28	1.28	0
2	10	20	1.92	3.20	1.92
3	10	30	2.88	6.08	2.88
4	10	30	4.32	9.12	3.04
5	10	30	6.48	13.68	4.56
6	10	30	9.72	20.52	6.85

Table 1: This table shows the number of racks purchased each year, the total number of racks, the amount of new storage purchased each year (in PB), and the net amount of new storage added each year (in PB). We assume a constant investment each year in new racks and an improvement in disk performance each year of 50%. We assume that after three years of use, we retire the racks.

	MapReduce	**Sphere UDF**
Storage	Disk data	Disk data and in-memory objects
Processing	Map and Reduce; Reduce can only exist after Map	Arbitrary user defined functions
Data exchange	Reducers “pull” result from previous mappers	UDFs “push” results to various bucket files
Data locality	Input data is assigned to the nearest Mapper (input locality)	Input data is assigned to the the nearest UDF (input locality). Output data can be sent to specific locations to coordinate with other processing (output locality).

Table 2: As the table shows, Hadoop’s implementation of MapReduce and Sphere’s implementation of user defined functions (UDF) support data locality in different ways.

4.2 Access to Bulk Downloads of the SDSS

For the past several years, the Open Cloud Consortium has provided infrastructure through Sector [20] so that the Sloan Digital Sky Survey (SDSS) [11] data can be downloaded in bulk. Currently, we make available SDSS data releases DR3, DR4, DR5, DR6, and DR7 for bulk download. These datasets are of sizes 2TB, 3TB, 4TB, 13TB, and 16TB, respectively.

These bulk SDSS downloads will be moved to the OSDC. In addition, we plan to make available some of the large datasets that are expected to be produced by the Milky Way Laboratory, which may be as large as 500 TB in size. The Milky Way Laboratory datasets will be able to be analyzed using the GrayWulf infrastructure [28].

We will be replicating the SDSS data in Hadoop, so that MapReduce [13] computations can be performed over the data within the OSDC. In addition, users can use the Sector/Sphere system [20] so that arbitrary User Defined Functions (UDF) can be invoked over the SDSS data and Milky Way Laboratory datasets.

5. ABOUT THE OCC

The Open Cloud Consortium (OCC) includes universities, such as the University of Illinois at Chicago, Northwestern University, the University of Chicago and Johns Hopkins University, companies, such as Yahoo! Inc. and Cisco, and government agencies, such as NASA. The goals of the OCC are: 1) to support the development of standards for cloud computing and frameworks for interoperating between clouds; 2) develop benchmarks for cloud computing; 3) support reference implementations for cloud computing, preferably open source reference implementations; 4) manage persistent cloud computing infrastructure, such as the Open Science Data Cloud; and 5) manage testbeds for cloud computing, such as the Open Cloud Testbed [17]. More information about the Open Cloud Consortium can be found on its web site [26].

We are in the process of setting up an advisory board for the OSDC, which among other responsibilities, will establish policies for selecting datasets for the OSDC.

6. RELATED WORK

The OSDC is similar to the Open Science Grid (OSG), Amazon’s public cloud, and the Internet Archive, but there are also important differences. As we describe these differences below, it is important to note that we are not making the case that the OSDC is better than these competing infrastructures but instead that these infrastructures are different and complementary.

The goal of the Open Science Grid software stack is “to provide a uniform computing and storage interface across many independently managed computing and storage clusters. Scientists, researchers, and students, organized as virtual organizations (VOs), are the consumers of the CPU cycles and storage [27].” There are two important differences between the Open Science Data Cloud (OSDC) and the Open Science Grid (OSG). First, the OSDC uses a hosted model in which one organization provides a managed resource to the community. In contrast, the OSG is model that uses grid services to federate multiple virtual organizations that are providing services. Second, the OSDC software stack is based upon cloud services (Eucalyptus, Hadoop and Sector), while the OSG software stack is based upon grid services (such as GRAM, GridFTP and Condor-G). Third, the goal of the OSDC is to support long term persistent access to scientific data, while the goal of the OSG is "to promote discovery and collaboration in data-intensive research by providing a computing facility and services that integrate distributed, reliable and shared resources to support computation at all scales [27]."

	Open Science Data Cloud	**Open Science Grid**
structure	organization	virtual organization
management	hosted	federated
key software	Hadoop, Eucalyptus, Sector	Globus, Condor
focus	data intensive computing	high performance computing
goal	OSDC will provide a long term persistent home for data and a cloud based infrastructure for analyzing it	OSG will enable VOs to use resources of other VOs

Although the Amazon public cloud includes support for Hadoop [3], there are still several challenges when working with large datasets using Amazon. First, for large datasets, the easiest way to get data into and out of the Amazon cloud, is to ship disks back and forth and to take advantage of a service that Amazon offers that will import and export data from these disks. Second, currently Amazon's S3 services limit the size of files to 5GB. Third, Amazon's S3 service does not currently support data locality. In contrast, the OSDC has high performance network connections to most universities and specialized software (UDT) that can exploit such networks. Second, there is no restriction within the OSDC on the size of files and the OSDC architecture can take advantage of data locality.

The Internet Archive is a 501(c)(3) not-for-profit that is designed to offer permanent access to researchers, historians, scholars and the general public a wide variety of material from the Internet, including documents, audio, and moving images [22]. In contrast, the OSDC is designed specifically to support scientific data, especially medium to large size scientific data. It is also designed to support in-place computing services over the data. Similar to the Internet Archive, it is organized as not-for-profit and has a goal of providing permanent storage for its digital assets.

7. STATUS

The Open Science Data Cloud is an outgrowth of the Open Cloud Testbed, which is an Open Cloud Consortium testbed for wide area large data clouds that can utilize 10 Gbps networks to connect geographically distributed data centers [17]. Both applications mentioned above in Section 4 currently use the Open Cloud Testbed.

We are currently in the process of installing 6 racks that will be dedicated exclusively to the OSDC. Once these racks are set up and available, the Open Cloud Testbed will be operated as an experimental infrastructure for software development, benchmarking and related activities, while the Open Science Data Cloud will be a persistent production resource for managing, analyzing, archiving and sharing scientific data.

Two of these racks will be placed at the University of Illinois at Chicago, two will be placed at Johns Hopkins University (for the SDSS and related projects), one will be placed at the University of Chicago/Argonne National Laboratory (for Cistrack and related projects), and one will be placed in a site that will be determined shortly. See Figure 3.

We have also begun a process to obtain stable longer term funding for the OSDC, which is necessary for the migration strategy described in Section 3.2 and so that the OSDC can commit to hosting additional datasets.

Acknowledgments

The Open Science Data Cloud is operated by the Open Cloud Consortium. The Open Cloud Consortium is managed by the Center for Computational Science Research, Inc., which is a 501(c)(3) not-for-profit corporation.

Equipment for the Open Science Data Cloud has been donated by Yahoo! Inc. This equipment will be used to host the Cistrack data, Sloan Digital Sky Survey data, and a variety of other datasets.

Wide area networking for the Open Science Data Cloud is provided through donations and partnerships with Cisco, the National Lambda Rail, and the StarLight communications exchange facility.

The development of the UDT and Sector software that is used by the OSDC was supported in part by the National Science Foundation through grants NSF grants OCI-0430781, CNS-0420847, and ACI-0325013.

The development of Cistrack was supported in part by by the Chicago Biomedical Consortium with support from The Searle Funds at The Chicago Community Trust. It was also supported by NIH through grants P50 GM081892 and U01 HG004264.

Figures 1 and 2 and the accompanying text are based in part on discussions with Malcolm Atkinson.

8. REFERENCES

[1] Amazon. Amazon Elastic Compute Cloud (amazon ec2). aws.amazon.com/ec2, 2008.

[2] Amazon. Amazon Web Services Developer Connection. Retrieved from http://aws.amazon.com, 2009.

[3] Amazon. Amazon Web Services. Retrieved from http://aws.amazon.com, 2010.

[4] M. Armbrust, A. Fox, R. Griffith, A. D. Joseph, R. H. Katz, A. Konwinski, G. Lee, D. A. Patterson, A. Rabkin, I. Stoica, , and M. Zaharia. Above the Clouds: A Berkeley View of Cloud Computing. Electrical Engineering and Computer Sciences,

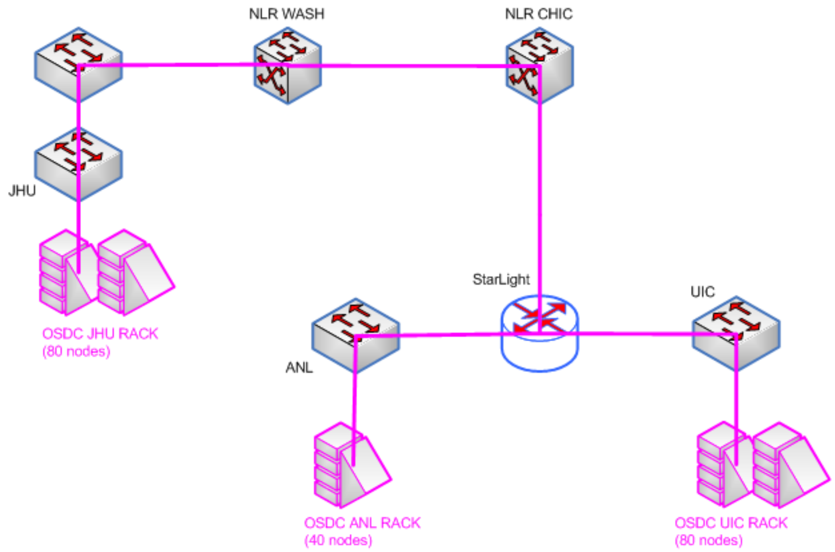

Figure 3: Initially, the hardware infrastructure for the Open Science Data Cloud will be distributed across three sites.

University of California at Berkeley, Technical Report No. UCB/EECS-2009-28.

[5] M. P. Atkinson, J. I. van Hemert, L. Han, A. Hume, and C. S. Liew. A distributed architecture for data mining and integration. In *DADC '09: Proceedings of the second international workshop on Data-aware distributed computing*, pages 11–20, New York, NY, USA, 2009. ACM.

[6] G. Bell, T. Hey, and A. Szalay. Beyond the data deluge. *Science*, 323(5919):1297 – 1298, 2009.

[7] K. D. Bollacker. Avoiding a digital dark age. *American Scientist*, 98:106–110, 2010.

[8] D. Borthaku. The Hadoop distributed file system: Architecture and design. retrieved from lucene.apache.org/hadoop, 2007.

[9] S. E. Celniker, L. A. L. Dillon, M. B. Gerstein, K. C. Gunsalus, S. Henikoff, G. H. Karpen, M. Kellis, E. C. Lai, J. D. Lieb, D. M. MacAlpine, G. Micklem, F. Piano, M. Snyder, L. Stein, K. P. White, R. H. Waterston, and modENCODE Consortium. Unlocking the secrets of the genome. *Nature*, 459(7249):927–930, June 2009.

[10] A. Chien, T. Faber, A. Falk, J. Bannister, R. Grossman, and J. Leigh. Transport protocols for high performance: Whither tcp? *Communications of the ACM*, 46(11):42–49, 2003.

[11] S. collaboration. Sloan Digitial Sky Survey. http://www.sdss.org, 2009.

[12] J. Dean and S. Ghemawat. MapReduce: Simplified data processing on large clusters. In *OSDI'04: Sixth Symposium on Operating System Design and Implementation*, 2004.

[13] J. Dean and S. Ghemawat. MapReduce: Simplified data processing on large clusters. *Communications of the ACM*, 51(1):107–113, 2008.

[14] Eucalyptus Systems. Eucalyptus - elastic utility computing architecture linking your programs to useful systems. http://www.eucalyptus.com/, 2009.

[15] I. Foster, Y. Zhao, I. Raicu, and S. Lu. Cloud computing and grid computing 360-degree compared. In *IEEE Grid Computing Environments (GCE08)*, pages 1–10, 2008.

[16] S. Ghemawat, H. Gobioff, and S.-T. Leung. The Google file system. In *SOSP '03: Proceedings of the nineteenth ACM symposium on Operating systems principles*, pages 29–43, New York, NY, USA, 2003. ACM.

[17] R. Grossman, Y. Gu, M. Sabala, C. Bennett, J. Seidman, and J. Mambratti. The open cloud testbed: A wide area testbed for cloud computing utilizing high performance network services. In *GridNets 2009*. Springer-Verlag, 2009.

[18] R. L. Grossman and Y. Gu. Data mining using high performance clouds: Experimental studies using Sector and Sphere. In *Proceedings of The 14th ACM SIGKDD International Conference on Knowledge Discovery and Data Mining (KDD 2008)*. ACM, 2008.

[19] Y. Gu and R. L. Grossman. UDT: UDP-based data transfer for high-speed wide area networks. *Computer Networks*, 51(7):1777—1799, 2007.

[20] Y. Gu and R. L. Grossman. Sector and sphere: Towards simplified storage and processing of large scale distributed data. *Philosophical Transactions of the Royal Society A, also arxiv:0809.1181*, 2009.

[21] Hadoop Project. Welcome to Apache Hadoop. retrieved from http://hadoop.apache.org, 2010.

[22] Internet Archive. About the Internet Archive, 2010.

[23] Large Hadron Collider. Worldwide lhc computing grid. http://lcg.web.cern.ch/lcg/.

[24] P. Mell and T. Grance. The nist definition of cloud computing, version 15. http://csrc.nist.gov/groups/SNS/cloud-computing/, October 15, 2009.

[25] C. Miceli, M. Miceli, S. Jha, H. Kaiser, and A. Merzky. Programming abstractions for data intensive computing on clouds and grids. In *9th IEEE/ACM*

International Symposium on Cluster Computing and the Grid (CCGrid 2009), pages 478–483, 2009.

[26] Open Cloud Consortium. http://www.opencloudconsortium.org, 2010.

[27] Open Science Grid. A national, distributed computing grid for data-intensive research. http://www.opensciencegrid.org, 2010.

[28] Y. Simmhan, R. Barga, C. van Ingen, M. Nieto-Santisteban, L. Dobos, N. Li, M. Shipway, A. S. Szalay, S. Werner, and J. Heasley. Graywulf: Scalable software architecture for data intensive computing. *Hawaii International Conference on System Sciences*, 0:1–10, 2009.

[29] StarLight. StarLight: The optical StarTap. http://www.startap.net/starlight, 2010.

[30] M. Stonebraker, J. Becla, D. Dewitt, K. Lim, D. Maier, O. Ratzesberger, , and S. Zdonik. Requirements for science data bases and scidb. In *Conference on Innovative Data Systems Research (CIDR)*. http://database.cs.brown.edu/projects/scidb/, 2009.

[31] A. S. Szalay, G. Bell, J. Vandenberg, A. Wonders, R. Burns, D. Fay, J. Heasley, T. Hey, M. Nieto-Santisteban, A. Thakar, C. van Ingen, and R. Wilton. Graywulf: Scalable clustered architecture for data intensive computing. *Hawaii International Conference on System Sciences*, 0:1–10, 2009.

[32] A. Thusoo, J. S. Sarma, N. Jain, Z. Shao, P. Chakka, S. Anthony, H. Liu, P. Wyckoff, and R. Murthy. Hive: a warehousing solution over a map-reduce framework. *Proceedings of VLDB*, 2(2):1626–1629, 2009.

[33] R. Wolski, C. Grzegorczyk, and D. N. et. al. Eucalyptus. retrieved from http://eucalyptus.cs.ucsb.edu/, 2008.