



HHS Public Access

Author manuscript

Trends Genet. Author manuscript; available in PMC 2020 March 01.

Published in final edited form as:

Trends Genet. 2019 March ; 35(3): 223–234. doi:10.1016/j.tig.2018.12.006.

Data Lakes, Clouds, and Commons: A Review of Platforms for Analyzing and Sharing Genomic Data

Robert L. Grossman^{1,*}

¹Center for Translational Data Science, University of Chicago, 900 East 57th Street, KCBD 10142, Chicago, IL 60637, USA

Abstract

Data commons collate data with cloud computing infrastructure and commonly used software services, tools, and applications to create biomedical resources for the large-scale management, analysis, harmonization, and sharing of biomedical data. Over the past few years, data commons have been used to analyze, harmonize, and share large-scale genomics datasets. Data ecosystems can be built by interoperating multiple data commons. It can be quite labor intensive to curate, import, and analyze the data in a data commons. Data lakes provide an alternative to data commons and simply provide access to data, with the data curation and analysis deferred until later and delegated to those that access the data. We review software platforms for managing, analyzing, and sharing genomic data, with an emphasis on data commons, but also cover data ecosystems and data lakes.

The Challenges of Large Genomic Datasets

The commoditization of sensors has resulted in new generations of instruments that produce large datasets that are available to genetics researchers. Next generation sequencing produced whole exome and whole genome datasets that were 200 to 800 GB or larger, and large projects such as The Cancer Genome Atlas (TCGA) [1] contain more than 2 PB of data and derived data.

Over the next few years, the research community will collect single-cell atlases [2], next generation imaging that captures the cellular microenvironment, and atlases about the cancer cells' interactions with the immunological system, all of which will produce ever larger datasets.

The accumulation of all these data has resulted in several challenges for the genetics research community. First, the size of the datasets is too large for all but the largest research organizations to manage and analyze. Second, the current model in which research groups set up their own computing infrastructure, download their own copy of the data, add their own data, and analyze the integrated dataset is simply too expensive for the government and private funding organizations to support. Third, the information technology (IT) expertise to

This is an open access article under the CC BY license (<http://creativecommons.org/licenses/by/4.0/>).

*Correspondence: robert.grossman@uchicago.edu (R.L. Grossman).

set up the required large-scale computing environments and the bioinformatics expertise to set up the required bioinformatics environments are difficult for most organizations to support. Fourth, because of **batch effects** (see Glossary) [3], it is usually considered wise to re-analyze all of the data (from raw data) using a common set of bioinformatics pipelines to minimize the presence of batch effects.

The importance of the appropriate data and computing infrastructure to create ‘knowledge bases’ and ‘knowledge networks’ to support precision medicine has been described in several reports [4,5].

In this review article, we describe some of the data, analysis, and collaboration platforms that have emerged to deal with these challenges.

Platforms for Data Sharing

Cloud Computing

Over the past 15 years, large-scale internet companies, such as Google, Amazon, and Facebook, have developed new computing infrastructure for their own internal use that became known as cloud computing platforms [6]. Some of these companies then made these platforms available to customers, including Amazon’s Amazon Web Services (AWS), Google’s Cloud Platform (GCP), and Microsoft’s Azure. Importantly, open source versions of some these platforms were also developed [7], including OpenStack (openstack.org) and OpenNebula (opennebula.org), enabling organizations to set up their own on-premise clouds. On-premise clouds are also called private clouds [8] to distinguish them from commercial public clouds that are used by multiple organizations.

NIST has developed a definition of cloud computing that includes the following characteristics [8]: (i) elastic in the sense that large-scale resources are available and (ii) self-provisioned in the sense that a user can provision the computing infrastructure required directly through a portal or **application programming interface (API)**.

Although it took time for cloud computing to be adapted for biomedical informatics, there was early recognition within the cancer community of the importance of this technology [9,10], and several university- and institute-based projects developed production-level cloud computing platforms to support the cancer research community, including the Bionimbus Protected Data Cloud [11], the Galaxy Cloud [12,13], Globus Genomics [14], and the Cancer Genome Collaboratory [15]. In addition, commercial companies, including DNAnexus [16] and Seven Bridges [17], developed cloud-based solutions for processing genomic data.

It may be helpful to divide computing platforms supporting biomedical research into three generations: (i) databases, (ii) **data clouds**, and (iii) **data commons** (Figure 1).

Databases and Data Portals

First generation platforms operated databases in which biomedical datasets were deposited, beginning with GenBank [18]. As the web became the dominant infrastructure for

collaboration, **data portals** emerged as applications that made the data in the underlying databases readily available to researchers. For the purposes here, one can think of a data portal as a website that provides interactive access to data in an underlying database. Although data portals are outside the scope of this review article, it is still important to mention the University of California, Santa Cruz (UCSC) Genome Browser [19] and cBioPortal [20] as some of the most important examples from this category.

The UCSC Genome Browser has been in continuous development since it was first launched in 2000 to help visualize the first working draft of the human genome assembly [19]. Today, it contains more than 160 assemblies from more than 90 species and can be run not only over the web but also downloaded and run locally using a version called genome browser in a box (GBiB) [21].

The cBioPortal for Cancer Genomics [20] is a widely used resource that integrates and visualizes cancer genomic data, including mutations, copy number variation, gene expression data, and clinical information. Currently, cBioPortal includes data from TCGA that are processed by Broad's Firehose and data from the International Cancer Genomic Consortium (ICGC) that are processed by the PANCAN Analysis Working group, plus additional smaller datasets [22]. cBioPortal was one of the first cancer data portals to organize data by genomic alterations, such as mutations, deletions, copy number variation, and expression levels, in a way that seemed natural to research oncologists and to tie the alterations back to the original cases to support further investigation when desired.

With next generation sequencing, the size of genomics datasets began to grow, and large-scale computing infrastructure is required to process, manage, and distribute data. Several systems were developed to process datasets such as the TCGA. CGHub was developed to host the **BAM** files from the TCGA project [23] by the UCSC. The Firehose system, developed by the Broad Institute, integrates data from TCGA and processes the data using applications from **Genome Analysis Toolkit (GATK)** [24] and applies algorithms such as **GISTIC2.0** [25] and **MutSig** [26]. The results can be browsed and accessed via a website (gdac.broadinstitute.org) and are available in Broad's FireCloud system [27].

Data Clouds

Second generation systems co-locate computing with biomedical data enabling researchers to compute over the data. A good example of this is the BLAST service [28] provided by the National Center for Biotechnology Information. Over the past decade, cloud computing has enabled the co-location of on-demand, large-scale computing infrastructure that has created new opportunities for the large-scale analysis of hosted biomedical data. Here, we use the term 'data cloud' for this integrated infrastructure. A working definition of a data cloud for biomedical data is a cloud computing platform [6] that manages and analyzes biomedical data and, usually, integrates the security and compliance required to work with controlled access biomedical data, such as germline genomic data. Examples of biomedical data clouds include the Bionimbus Protected Data Cloud developed by the University of Chicago [11], the Cancer Genomics Cloud developed by Seven Bridges Genomics [17], the Cancer Collaboratory developed by the Ontario Institute for Cancer Research [29], and the Galaxy Cloud [12,13] developed by the Galaxy Project.

We now describe three important milestones in the use of large-scale cloud computing in genomics. The first milestone was the launch of the National Cancer Institute (NCI) Genomics Data Commons [30] that used an OpenStack-based private cloud to analyze and harmonize genomic and associated clinical data from more than 18 000 cancer-tumor-normal pairs, including TCGA [1]. By **data harmonization**, we mean applying a uniform set of pipelines for cleaning, applying quality control criteria, processing, and post-processing submitted data[31]. The second milestone was the development of the three NCI Cloud Pilots: the ISB Cancer Genomics Cloud by the Institute for Systems Biology [32], FireCloud by the Broad Institute [27], and the Cancer Genomics Cloud by Seven Bridges Genomics [17], each of which provided cloud-based computing infrastructure to analyze TCGA data. The first two Cloud Pilots used GCP and the third used AWS. A third important milestone was the analysis of 280 whole genomes using multiple distributed public and private clouds by the PANCAN Analysis Working Group [29].

Cloud computing is widely used today to support scientific research for many disciplines outside of the biomedical sciences. In general, the architecture for these systems is simpler since the security and compliance infrastructure required for working with controlled access biomedical data is not required.

Data Commons

Third generation systems integrate biomedical data, computing and storage infrastructure, and software services required for working with data to create a data commons. Some examples of data commons and six core requirements for data commons are reviewed in [33]. A working definition of a ‘data commons’ is the colocation of data with cloud computing infrastructure and commonly used software services, tools, and applications for managing, integrating, analyzing, and sharing data that are exposed through APIs to create an interoperable resource [33].

Some of the core services (data common services) required for a data commons are as follows:

- i.** authentication services for identifying researchers;
- ii.** authorization services for determining which datasets researchers can access;
- iii.** digital ID services for assigning permanent identifiers to datasets and accessing data using these IDs;
- iv.** metadata services for assigning metadata to a digital object identified by a digital ID and accessing the metadata;
- v.** security and compliance services so that data commons can support controlled access data;
- vi.** data model services for integrating data with respect to one or more data models; and
- vii.** workflow services for executing bioinformatics pipelines so that data can be analyzed and harmonized.

Accessing controlled access data requires services (i) and (ii). With service (iii), data stored in commons are findable and accessible. With service (iv), data stored in data commons can be reusable and interoperable. In practice, for data to be reusable depends in large part on the quality of the data annotation prepared by the data submitter. With services (iii) and (iv), data stored in commons are findable, accessible, reusable, and interoperable and thus is sometimes abbreviated as FAIR. The importance of making biomedical data FAIR has been stressed in efforts such as the European FORCE11 Initiative [34] and the National Institutes of Health (NIH) Big Data to Knowledge (BD2K) initiative [35]. Recently, a framework for metrics to measure the ‘FAIRness’ of services has also been developed [36].

Workflow services (vii) in data commons are quite varied and include running existing workflows that have been integrated into the commons and can be used to analyze data in the commons, pulling existing workflows from workflow repositories outside the commons and applying them to data in the commons, and developing new workflows and using them to analyze data in the commons. Also, some commons allow users to execute workflows, while others limit this to the data commons administrators.

An example of a data commons is the NCI Genomic Data Commons (GDC) [30,37], used by more than 100 000 distinct cancer researchers in 2018. With the GDC [30], data commons began to curate and integrate contributed data using a common data model [core service (vi)], harmonize the contributed data using a common set of bioinformatics pipelines [core service (vii)], support the visual exploration of data through a data portal, and expose APIs to the core services (i)-(v) to support third party applications over the integrated and harmonized data.

Project Data: Object Data and Structured Data

Data in a data commons are usually organized into projects, with different projects potentially having their own data model and collecting different subsets of clinical, molecular, imaging, and other data. It is an open question of how best to organize data across projects so that it can be integrated, harmonized, and queried. One natural division that is emerging is the distinction between the **structured data**, the unstructured data, and the **data objects** in a project. The object data typically include FASTQ or BAM files [38] used in genomics, image files, video files, and other large files, such as archive or backup files associated with a project. The structured data include clinical data, demographic data, biospecimen data, variant data [39], and other data associated with a data schema. The unstructured data include text, notes, articles, and other data that are not associated with a schema.

Part of the curation process is to align the structured data in a project with an appropriate ontology. Examples include using the human phenotype ontology [40] and the NCI Thesaurus [41] for curating clinical phenotype data and CDISC [42] for curating clinical trials data. It can be quite challenging and labor intensive to match ontologies to clinical data, and several tools have been developed to make this easier [43,44].

If we call all the structured data, unstructured data, and associated schemas ‘project core data’, then it is quite common for the project’s object data to be 1000 times (or more) larger

than the project's core data. For example, with the TCGA's projects [1], the data objects were measured in 10s to 100s TB, while the project's core data were measured in 10s of GB.

In practice, a project's object data are assigned **globally unique identifiers (GUIDs)** and metadata and stored in clouds using services (iii) and (iv) and are immutable (although new versions may be added to the project), while the project core data are often updated, as part of the curation and quality assurance process and as new data are added to the project.

A project's object data are searched via its metadata [core service (iv)], while a project's core data can be searched via its data model [core service (vi)]. Of course, a project's object data can be processed to produce features that can then be managed and searched. Examples include developing algorithms for identifying particular types of cells in cell images and searching for these cells or processing BAM files to compute data quality scores and searching for BAM files with particular data quality problems. When data are curated and integrated with a common data model, synthetic cohorts can be created through a query, such as 'find all males over 50 years of old that smoked and have a KRAS mutation [45].'

Another way to think of this is that core services (i)-(v) support the 'shallow' indexing and search via metadata, while core services (i)-(vi) support 'deep' indexing and search via the data model attached to project core data. In either case, when the services are exposed via APIs to third party applications, data become portable and data commons become interoperable, both of which are usually thought of as important requirements [33].

Data Lakes

Sometimes the term '**data lake**' is used when data are stored simply with digital IDs and metadata (shallow indexing), but without a data model. Data models and schemas are used when the data are written or when the data are analyzed, but not when the data are stored. Additional information about data lakes can be found in [46]. Since it can be very labor intensive to import data with respect to a data model, and since not all the data in a commons are used, this has the advantage that the effort to align the data with a data model is not needed until the data are analyzed. Of course, at the time the data are analyzed and aligned with a data model, the expertise to do this may no longer be easily available.

Through the use of cloud computing, data commons can support large-scale data, but this also creates sustainability challenges, due to the cost of large-scale storage and compute. One sustainability model that can be attractive to an organization is to provide the data at no cost, but to control the cost of the computing resources by using a 'pay for compute model' [33], establishing quotas for compute, giving compute allocations, or distributing 'chits' that can be redeemed for compute.

Just as data lakes required less curation than data commons, data catalogs required less curation than data lakes. A 'data catalog' is simply a listing of data assets, some basic metadata, and their locations, but without a common mechanism for accessing the data, such as used in a data lake.

Workflows

Bioinformatics workflows are often data intensive and complex, consisting of several different programs with the outputs of one program used as the inputs to another. For this reason, specialized workflow management systems have been developed so that workflows can be mapped efficiently to different high-performance, parallel, and distributed computer systems. Workflow languages have been developed so that domain specialists knowledgeable about the workflows can describe the workflows in a manner that is independent of the specific underlying physical architecture of the system executing the workflows. Despite many years of effort though, there is still no standard language for expressing workflows in general and bioinformatics workflows in particular [47,48]. Within the cancer genomics community, the Common Workflow Language (CWL) [49] is gaining in popularity. The GA4GH Consortium (ga4gh.org) supports a technical effort to standardize bioinformatics workflows, which includes the workflow execution services (WES) and task execution service (TES). With the growing use of **container**-based environments for program execution, such as **Docker**, it is becoming more common to encapsulate workflows in containers to make them easier to reuse [50]. Before the wide adoption of containers, workflows were encapsulated in virtual machines for the same reason. Examples of services for accessing reproducible workflows include Dock-store [51] and Biocompute Objects [52].

Data and Commons Governance

A common definition of IT governance is [53]: (i) Assure that the investments in IT generate business value. (ii) Mitigate the risks that are associated with IT. (iii) Operate in such a way as to make good long-term decisions with accountability and traceability to those funding IT resources, those developing and support IT resources, and those using IT resources. This definition can be easily adapted to provide a good definition for data commons governance: (i) Assure that the investments in the data commons generate value to the research community. (ii) Manage the balance between the risks associated with participant data and the benefits realized from research involving these data [54]. (iii) Operate in such a way so as to make good long-term decisions with accountability and traceability to those sponsors that fund the data commons; the engineers that develop, manage, and operate the data commons; and the researchers that use it.

An overview of principles for data commons and a description of eight principles for biomedical data commons can be found in [55]. A survey of how data are made available and controlled in commons is in [56]. A survey of data commons governance models is in [57]. The GA4GH framework for sharing data is described in [54].

The data governance structure for international data commons, such as the INRG Data Commons [58] and the ICGC Data Commons [59], can be challenging and may have restrictions on the movement of the underlying controlled access genomic data.

Building and Operating a Data Commons

Building a data commons usually consists of the following steps (Figure 2):

- i. Put in place data governance agreements that govern the contribution, management and use of the data in the data commons and common governance agreements that govern the development, operations, use, and sustainability of the commons.
- ii. Develop a data model (or data models) that describe the data in the commons.
- iii. Set up and configure the data commons itself.
- iv. Work with the community to submit data to the data commons.
- v. Import, clean, and curate the submitted data.
- vi. Process and analyze the data using bioinformatics pipelines to produce harmonized data products. This is often done with analysis working groups.
- vii. Open up the commons to external researchers, third party applications, and interoperate with other commons.

To support the activities, a data commons usually has the following components:

- i. A data exploration portal (or more simply a data portal) for viewing, exploring, visualizing and downloading the data in the commons.
- ii. A data submission portal for submitting data to the commons.
- iii. An API supporting third party applications.
- iv. Systems for the large-scale processing of data in the commons to produce derived data products.
- v. Systems to support analysis working groups and other team science constructs used for the collaborative analysis and annotation of data in the commons. What are being called ‘workspaces’ are one of the mechanisms that are emerging to support this.

Data Ecosystems Containing Multiple Data Commons

As the number of data commons grow, there will be an increasing need for data commons to interoperate and for applications to be able to access data and services from multiple data commons. It may be helpful to think of this situation as laying the foundation for a data ecosystem [60].

Sometimes the data commons services (i)-(vi) described above are called ‘framework services’ since they provide the framework for building a data commons, and, in fact, can be used to support multiple data commons that interoperate (Figure 3, Key Figure). As mentioned above, when these services are exposed through an API, either as part of a data commons or as a part of framework services supporting multiple data commons, they can support an ecosystem of third party applications [45].

There is no generally accepted definition of a data ecosystem at this time, but, at the minimum, a ‘*data ecosystem*’ for biomedical data (as opposed to a data commons) should support the following:

- i. Authentication and authorization services so that a community of researchers can access an ecosystem of data and applications with a common (research) identity and common authorization that is shared across data commons and applications.
- ii. A collection of applications that are powered by APIs that are FAIR compliant that are shared across multiple data commons.
- iii. The ability for multiple data commons to interoperate through framework services and, preferably, through data peering [33] so that access to data across data commons and applications is transparent, frictionless, and without egress charges, as long as the access is through a digital ID.
- iv. Shared data models, or portions of data models, to simplify the ability for third party applications to access data from multiple data commons and applications. Projects within a larger overall program, or in related programs, may share a data model. More commonly, different projects may share some common data elements within a core data model, with each project having additional data elements unique for that project.
- v. Support for workspaces that may include the following:
 - a. the ability to create synthetic (or virtual) cohorts and export cohorts to workspaces;
 - b. the ability to execute bioinformatics workflows within workspaces; and
 - c. workspace services for processing, exploring, and analyzing data using containers, virtual machines, or other mechanisms.
- vi. Security and compliance services.

Often workspace services (vb) and (vc) use a user-pay model as mentioned above.

An example of a cancer data ecosystem is the NCI Cancer Research Data Commons or NCRDC [61]. The NCRDC spans the GDC [30] and the Cloud Resources [61], so that both AWS and the GCP can be used to both analyze data from the GDC as well as to support integrative data analysis across data uploaded by researchers with data from the GDC and other third party datasets. Data commons for proteomic and imaging data are in the process of being added to the NCRDC. The NCRDC uses the framework services described above so that multiple data commons and other NCRDC resources can share authentication, authorization, ID, and metadata services. In particular, this approach allows applications to be built that span multiple data commons.

Concluding Remarks and Future Directions

We have reviewed some of the more recent data and computing platforms that have been used to analyze large-scale data being produce in biology, medicine, and health care, with a particular emphasis on data commons. See Figure 4 for an overview of the different platforms. Data commons provide several important advantages, including the following:

- Data commons support repeatable, reproducible, and open research.

- Some diseases are dependent upon having a critical mass of data to provide the required statistical power for the scientific evidence (e.g., to study combinations of rare mutations in cancer).
- With more data, smaller effects can be studied (e.g., to understand the effect of environmental factors on disease).
- Data commons enable researchers to work with large datasets at much lower cost to the sponsor than if each researcher set up their own local environment.
- Data commons generally provide higher security and greater compliance than most local computing environments.
- Data commons support large-scale computation so that the latest bioinformatics pipelines can be run.
- Data commons can interoperate with each other so that over time data sharing can benefit from a ‘network effect’.

Over the next few years, one of the most important changes will be the ability of patients to submit their own data to a data commons and to gain some understanding of their own data in terms of the overall data available in the commons and their broader data ecosystem the commons is part of (see Outstanding Questions). The ability of patients to contribute their own data and to have control over how the data are used by the research community [62] is an important aspect of what is sometimes called patient partnered research.

Acknowledgments

This project has been funded in part with federal funds from the NCI, NIH, task order 17X053 and task order 14X050 under contract HHSN261200800001E. The content of this publication does not necessarily reflect the views or policies of the Department of Health and Human Services, nor does mention of trade names, commercial products, or organizations imply endorsement by the US Government.

Glossary

Application programming interface (API)

An API is a specification for how two different software programmers communicate with each other and an implementation of the specification in computer code

BAM

The binary alignment map is a binary format that is widely used for storing molecular sequence data

Batch effects

Batch effects are differences in samples that are the results of differences in laboratory conditions; materials used to prepare the samples, such as reagents; personnel that prepare the samples; and other differences. Batch effects are often an important confounding factor in high-throughput sequence data

Container

A container for running software is a package of software that includes everything needed to run a software application, including the application's code, as well as the runtime environment, system tools, system libraries, configurations, and settings. Containers are designed to be run in different types of computing environments with no changes

Data cloud

A data cloud is a cloud computing platform for managing, analyzing, and sharing datasets

Data commons

A data commons co-locates data with cloud computing infrastructure and commonly used software services, tools, and applications for managing, integrating, analyzing, and sharing data that are exposed through APIs to create an interoperable resource

Data harmonization

Data harmonization is the process that brings together data from multiple sources and applies uniform and consistent processes, such as uniform quality control metrics to the accepted data; mapping the data to a common data model; processing the data with common bioinformatics pipelines; and post-processing the data using common quality control metrics

Data lake

A data lake is a system for storing data as objects, where the objects have an associated GUID and (object) metadata, but there is no data model for interpreting the data within the object

Data object

In cloud computing, a data object consists of data, a key, and associated metadata. The data can be retrieved using key and the metadata associated with a specific data object can be retrieved, but more general queries are not supported. Amazon's S3 storage system is a widely used storage system for data objects

Data portal

A data portal is a website that provides interactive access to data in an underlying data management system, such as a database. Data commons and data lakes can also have data portals

Docker

Docker is a software program for running containers developed by the company Docker, Inc. The containers it runs are often called Docker containers

Genome Analysis Toolkit (GATK)

GATK is a widely used collection of bioinformatics pipelines and associated best practices for variant discovery and genotyping developed by the Broad Institute

GISTIC2.0

GISTIC is a probabilistic algorithm for detecting somatic copy number alterations that are likely to drive cancer growth

Globally Unique Identifier (GUID)

A GUID is an essentially unique identifier that is generated by an algorithm so that no central authority is needed, but rather different programs running in different locations can generate GUID with a low probability that they will collide. A common format for a GUID is the hexadecimal representation of a 128-bit binary number

MutSig

MutSig (for mutation significance) is a probabilistic algorithm and associated software application that analyzes a list of mutations produced from DNA sequencing data to identify genes that were mutated more often than expected by chance, given background mutation processes

NIST

The National Institute of Standards and Technology is a US federal agency that advances measurement science and develops standards. NIST has developed definitions in standards for cloud computing and information security

Structured data

Data are structured if they are organized into records and fields, with each record consisting of one or more data elements (data fields). In biomedical data, data fields are often restricted to controlled vocabularies to make querying them easier

References

1. Tomczak K et al. (2015) The Cancer Genome Atlas (TCGA): an immeasurable source of knowledge. *Contemp. Oncol. (Pozn.)* 19, A68–A77 [PubMed: 25691825]
2. Rozenblatt-Rosen O et al. (2017) The Human Cell Atlas: from vision to reality. *Nature* 550, 451–453 [PubMed: 29072289]
3. Leek JT et al. (2010) Tackling the widespread and critical impact of batch effects in high-throughput data. *Nat. Rev. Genet* 11, 733–739 [PubMed: 20838408]
4. Council NR (2011) *Toward Precision Medicine: Building a Knowledge Network for Biomedical Research and a New Taxonomy of Disease*, The National Academies Press
5. Panel BR (2016) Cancer Moonshot Blue Ribbon Panel Report. <https://www.cancer.gov/research/key-initiatives/moonshot-cancer-initiative> (accessed 2018)
6. Armbrust M et al. (2009) Above the clouds: a Berkeley view of cloud computing. Technical report UCB/EECS-2009–28, EECS Department, University of California, Berkeley
7. Von Laszewski G. Comparison of multiple cloud frameworks. *Cloud Computing (CLOUD)*, 2012 IEEE 5th International Conference on, IEEE; 2012. 734–741.
8. Mell P and Grance T (2011) The NIST definition of cloud computing (draft): recommendations of the National Institute of Standards and Technology, National Institute of Standards and Technology
9. Grossman RL and White KP (2012) A vision for a biomedical cloud. *J. Intern. Med* 271, 122–130 [PubMed: 22142244]
10. Stein LD (2010) The case for cloud computing in genome informatics. *Genome Biol.* 11, 207 [PubMed: 20441614]
11. Heath AP et al. (2014) Bionimbus: a cloud for managing, analyzing and sharing large genomics datasets. *J. Am. Med. Inform. Assoc* 21, 969–975 [PubMed: 24464852]
12. Afgan E et al. (2010) Galaxy CloudMan: delivering cloud compute clusters. *BMC Bioinform.* 11, S4
13. Afgan E et al. (2011) Harnessing cloud computing with Galaxy Cloud. *Nat. Biotechnol* 29, 972–974 [PubMed: 22068528]

14. Madduri RK et al. (2014) Experiences building Globus Genomics: a next-generation sequencing analysis service using Galaxy, Globus, and Amazon Web Services. *Concurr. Comput* 26, 2266–2279 [PubMed: 25342933]
15. Yung CK et al. (2017) The Cancer Genome Collaboratory, abstract 378, AACR
16. Shringarpure SSet et al. (2015) Inexpensive and highly reproducible cloud-based variant calling of 2,535 human genomes. *PLoS One* 10, e0129277 [PubMed: 26110529]
17. Lau JW et al. (2017) The Cancer Genomics Cloud: collaborative, reproducible, and democratized—a new paradigm in large-scale computational research. *Cancer Res.* 77, e3–e6 [PubMed: 29092927]
18. Benson D et al. (1993) GenBank. *Nucleic Acids Res.* 21, 2963–2965
19. Kent WJ et al. (2002) The human genome browser at UCSC. *Genome Res.* 12, 996–1006 [PubMed: 12045153]
20. Gao J et al. (2013) Integrative analysis of complex cancer genomics and clinical profiles using the cBioPortal. *Sci. Signal* 6, pl 1
21. Rosenbloom KR et al. (2015) The UCSC Genome Browser database: 2015 update. *Nucleic Acids Res.* 43, D670–D681 [PubMed: 25428374]
22. Cerami E et al. (2012) The Cancer Genomics portal: an open platform for exploring multidimensional cancer genomics data. *Cancer Discov.* 2, 401–404 [PubMed: 22588877]
23. Wilks C et al. (2014) The Cancer Genomics Hub (CGHub): overcoming cancer through the power of torrential data. *Database (Oxford)* 2014, bau093 [PubMed: 25267794]
24. DePristo MA et al. (2011) A framework for variation discovery and genotyping using next-generation DNA sequencing data. *Nat. Genet* 43, 491 [PubMed: 21478889]
25. Mermel CH et al. (2011) GISTIC2.0 facilitates sensitive and confident localization of the targets of focal somatic copy-number alteration in human cancers. *Genome Biol.* 12, R41 [PubMed: 21527027]
26. Lawrence MS et al. (2013) Mutational heterogeneity in cancer and the search for new cancer-associated genes. *Nature* 499, 214 [PubMed: 23770567]
27. Birger C et al. (2017) FireCloud, a scalable cloud-based platform for collaborative genome analysis: strategies for reducing and controlling costs. *bioRxiv* 209494
28. Boratyn GM et al. (2013) BLAST: a more efficient report with usability improvements. *Nucleic Acids Res.* 41, W29–W33 [PubMed: 23609542]
29. Yung CK et al. (2017) Large-scale uniform analysis of cancer whole genomes in multiple computing environments. *bioRxiv* Published online 7 10, 2017. 10.1101/161638
30. Grossman RL et al. (2016) Toward a shared vision for cancer genomic data. *N. Engl. J. Med* 375, 1109–1112 [PubMed: 27653561]
31. Lee JS-H et al. (2018) Data harmonization for a molecularly driven health system. *Cell* 174, 1045–1048 [PubMed: 30142341]
32. Reynolds SM et al. (2017) The ISB Cancer Genomics Cloud: a flexible cloud-based platform for cancer genomics research. *Cancer Res.* 77, e7–e10 [PubMed: 29092928]
33. Grossman RL et al. (2016) A case for data commons: toward data science as a service. *Comput Sci. Eng* 18, 10–20 [PubMed: 29033693]
34. Wilkinson MD et al. (2016) The FAIR Guiding Principles for scientific data management and stewardship. *Sci. Data* 3, 160018 [PubMed: 26978244]
35. Bourne PE et al. (2015) The NIH Big Data to Knowledge (BD2K) initiative. *J. Am. Med. Inform. Assoc* 22, 1114 [PubMed: 26555016]
36. Wilkinson MD et al. (2017) A design framework and exemplar metrics for FAIRness. *bioRxiv* 225490
37. Jensen MA et al. (2017) The NCI Genomic Data Commons as an engine for precision medicine. *Blood* 130, 453–459 [PubMed: 28600341]
38. Clarke L et al. (2012) The 1000 Genomes Project: data management and community access. *Nat. Methods* 9, 459 [PubMed: 22543379]
39. Danecek P et al. (2011) The variant call format and VCFtools. *Bioinformatics* 27, 2156–2158 [PubMed: 21653522]

40. Köhler S et al. (2016) The human phenotype ontology in 2017. *Nucleic Acids Res.* 45, D865–D876 [PubMed: 27899602]
41. Sioutos N et al. (2007) NCI Thesaurus: a semantic model integrating cancer-related clinical and molecular information. *J. Biomed. Inform* 40, 30–43 [PubMed: 16697710]
42. Huser V et al. (2015) Standardizing data exchange for clinical research protocols and case report forms: an assessment of the suitability of the Clinical Data Interchange Standards Consortium (CDISC) Operational Data Model (ODM). *J. Biomed. Inform* 57, 88–99 [PubMed: 26188274]
43. Mungall CJ et al. (2016) The Monarch Initiative: an integrative data and analytic platform connecting phenotypes to genotypes across species. *Nucleic Acids Res.* 45, D712–D722 [PubMed: 27899636]
44. Haendel M et al. (2018) A census of disease ontologies. *Annu. Rev. Biomed. Data Sci* 1, 305–331
45. Wilson S et al. (2017) Developing cancer informatics applications and tools using the NCI Genomic Data Commons API. *Cancer Res.* 77, e15–e18 [PubMed: 29092930]
46. Terrizzano IG et al. (2015) Data wrangling: the challenging journey from the wild to the lake, CIDR
47. Leipzig J (2017) A review of bioinformatic pipeline frameworks. *Brief. Bioinform* 18, 530–536 [PubMed: 27013646]
48. Alterovitz G et al. (2018) Enabling precision medicine via standard communication of NGS provenance, analysis, and results. *bioRxiv* 191783
49. Amstutz P et al. (2016) Common Workflow Language, v1.0. Specification, Common Workflow Language working group, <https://www.commonwl.org/>
50. Boettiger C (2015) An introduction to Docker for reproducible research. *ACM SIGOPS Oper. Syst Rev* 49, 71–79
51. O'Connor BD et al. (2017) The Dockstore: enabling modular, community-focused sharing of Docker-based genomics tools and workflows. *F1000Research* 6
52. Simonyan V et al. (2017) Biocompute objects—a step towards evaluation and validation of biomedical scientific computations. *PDA J. Pharm. Sci. Technol* 71, 136 [PubMed: 27974626]
53. Brown AE and Grant GG (2005) Framing the frameworks: a review of IT governance research. *Commun. Assoc. Inf. Syst* 15, 38
54. Knoppers BM (2014) Framework for responsible sharing of genomic and health-related data. *Hugo J.* 8, 3 [PubMed: 27090251]
55. Deverka PA et al. (2017) Creating a data resource: what will it take to build a medical information commons? *Genome Med.* 9, 84 [PubMed: 28938910]
56. Eschenfelder KR and Johnson A (2014) Managing the data commons: controlled sharing of scholarly data. *J. Assoc. Inf. Sci. Technol* 65, 1757–1774
57. Fisher JB and Fortmann L (2010) Governing the data commons: policy, practice, and the advancement of science. *Inf. Manag* 47, 237–245
58. Volchenboum SL et al. (2017) Data commons to support pediatric cancer research In *American Society of Clinical Oncology Educational Book. American Society of Clinical Oncology Meeting, 2017*, pp. 746–752
59. Zhang J et al. (2011) International Cancer Genome Consortium Data Portal—a one-stop shop for cancer genomics data. *Database (Oxford)* bar026 [PubMed: 21930502]
60. Grossman RL (2018) Progress toward cancer data ecosystems. *Cancer J.* 24, 122–126
61. Hinkson IV et al. (2017) A comprehensive infrastructure for big data in cancer research: accelerating cancer research and precision medicine. *Front. Cell Dev. Biol* 5, 83 [PubMed: 28983483]
62. Wilbanks J and Friend SH (2016) First, design for data sharing. *Nat. Biotechnol* 34, 377 [PubMed: 26939011]

Highlights

Data commons collate data with cloud computing infrastructure and commonly used software services, tools, and applications to create biomedical resources for the large-scale management, analysis, harmonization, and sharing of biomedical data.

Data commons support repeatable, reproducible, and open research.

Data lakes provide access to a collection of data objects that can be accessed via digital IDs and searched via their metadata.

A simple data ecosystem can be built when a data commons exposes an API that can support a collection of third party applications that can access data from the commons. More complex data ecosystems arise when multiple data commons can interoperate and support a collection of third party applications over a common set of core services (framework services), such as services for authentication, authorization, digital IDs, and metadata.

Outstanding Questions

In practice, uploading clinical phenotype data into a data commons so that it is aligned with the data common's data model and can be harmonized with the other data in the commons is quite labor intensive. An open question is how to develop bioinformatics tools and associated frameworks so that data can be transformed automatically or semi-automatically into the proper format.

Developing software architectures and associated platforms that can that can query and aggregate data from multiple data commons is an important challenge.

In general, different commons will have both large and small differences in the practices and standards used for assigning clinical phenotype. Developing applications that can query and aggregate data from multiple data commons even when there are minor (or major) differences between the clinical phenotype data is an important challenge.

In practice, researchers will be analyzing data using applications that are hosted across multiple commercial public clouds, while those operating data commons will try to reduce their costs by focusing on one or two public or private clouds. What are the software architectures and operating procedures so that data commons can operate across just one or two public or private clouds but support researchers across multiple clouds?

Moving data projects between data commons is important so that data commons do not begin to 'silo' data. What are appropriate serialization for-mats so that projects can be efficiently imported and exported between data commons?



Figure 1. Some of the Important Differences between Data Clouds and Data Commons.

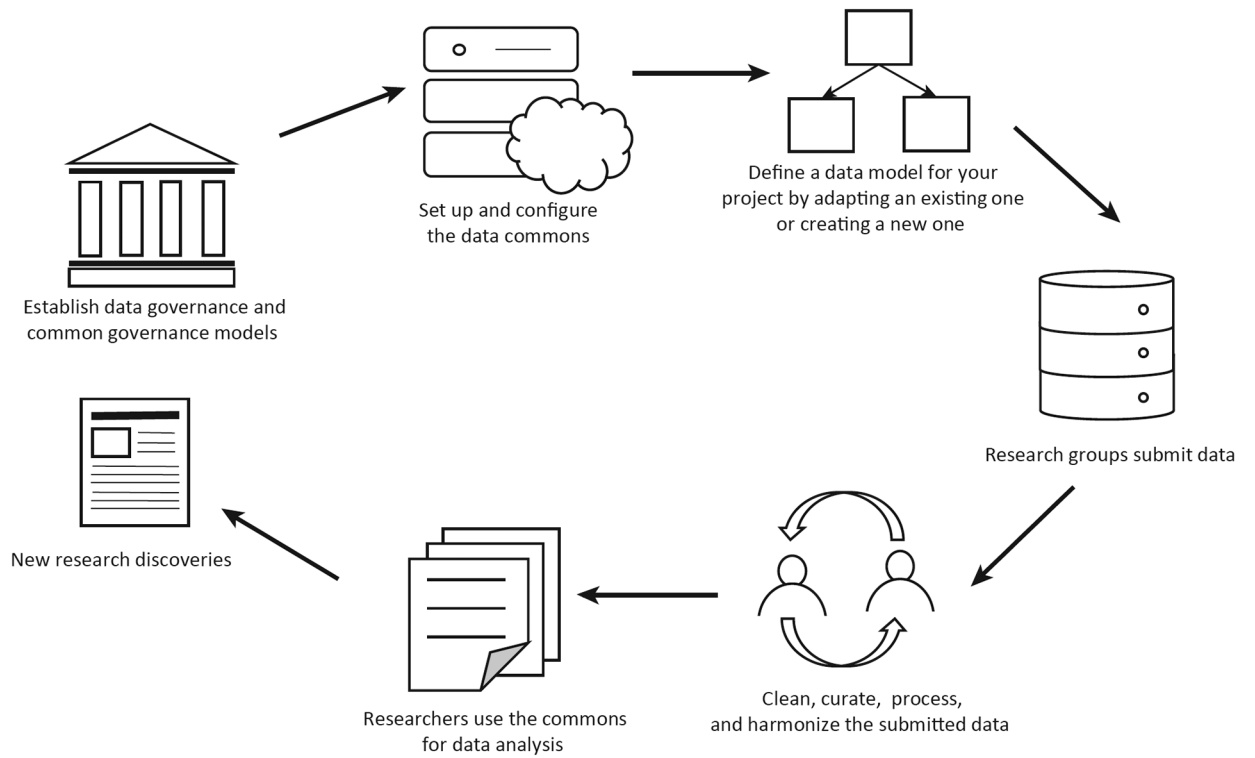


Figure 2. Building a Data Commons.

Data commons support the entire life cycle of data, including defining the data model, importing data, cleaning data, exploring data, analyzing data, and then sharing new research discoveries.

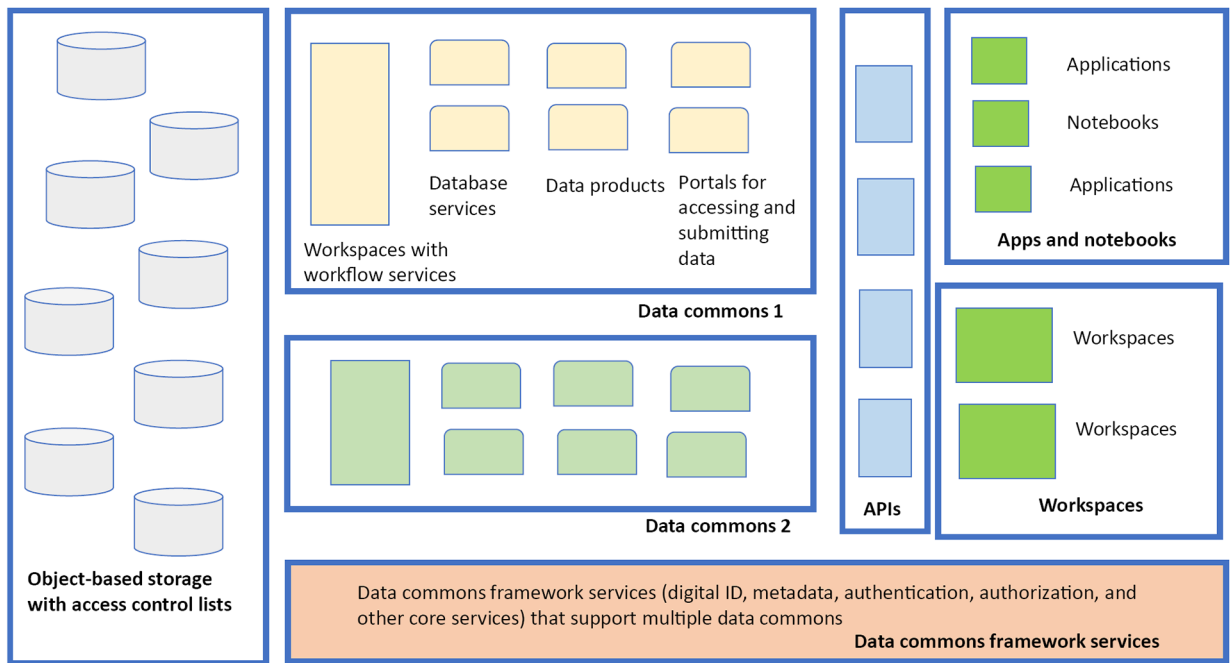


Figure 3. Key Figure Data Commons Framework Services

This diagram shows how data commons framework services can support multiple data commons and an ecosystem of workspaces, notebooks, and applications.

Author Manuscript

Author Manuscript

Author Manuscript

Author Manuscript

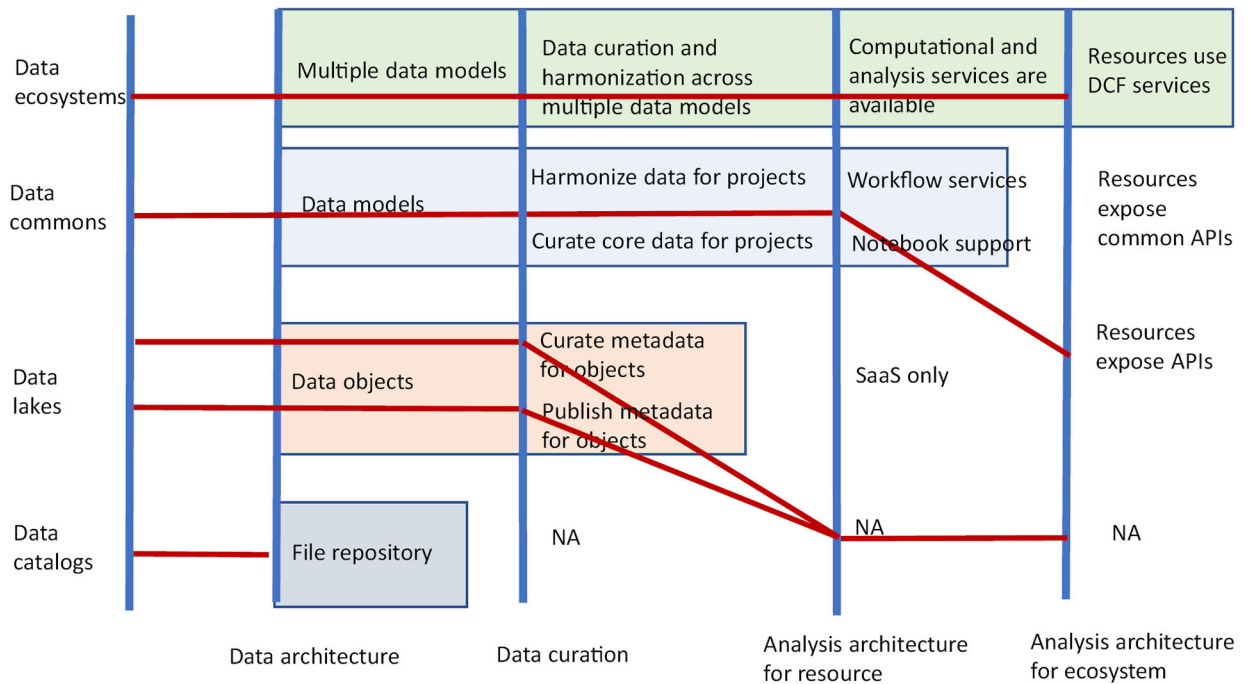


Figure 4. Data Platforms.

Data platforms can be categorized along four axes: the data architecture, the extent of the data curation and harmonization, the analysis architecture of a resource, and the analysis architecture of the ecosystem. The red lines can be viewed as classifying platforms using parallel coordinates and these four dimensions. The top line is the parallel coordinates associated with the National Cancer Institute (NCI) Cancer Research Data Commons, the line below is the parallel coordinates for the NCI Genomic Data Commons, the two lines below are two possible architectures for data lakes, while the bottom line is an architecture for a repository of files. Abbreviations: API, application programming interface; DCF, data commons framework; NA, not applicable; SaaS, software as a service.