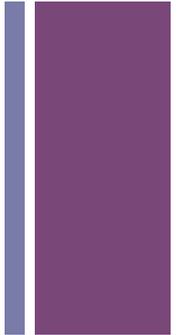


# Big Data!

Saba Sehrish

# + Outline

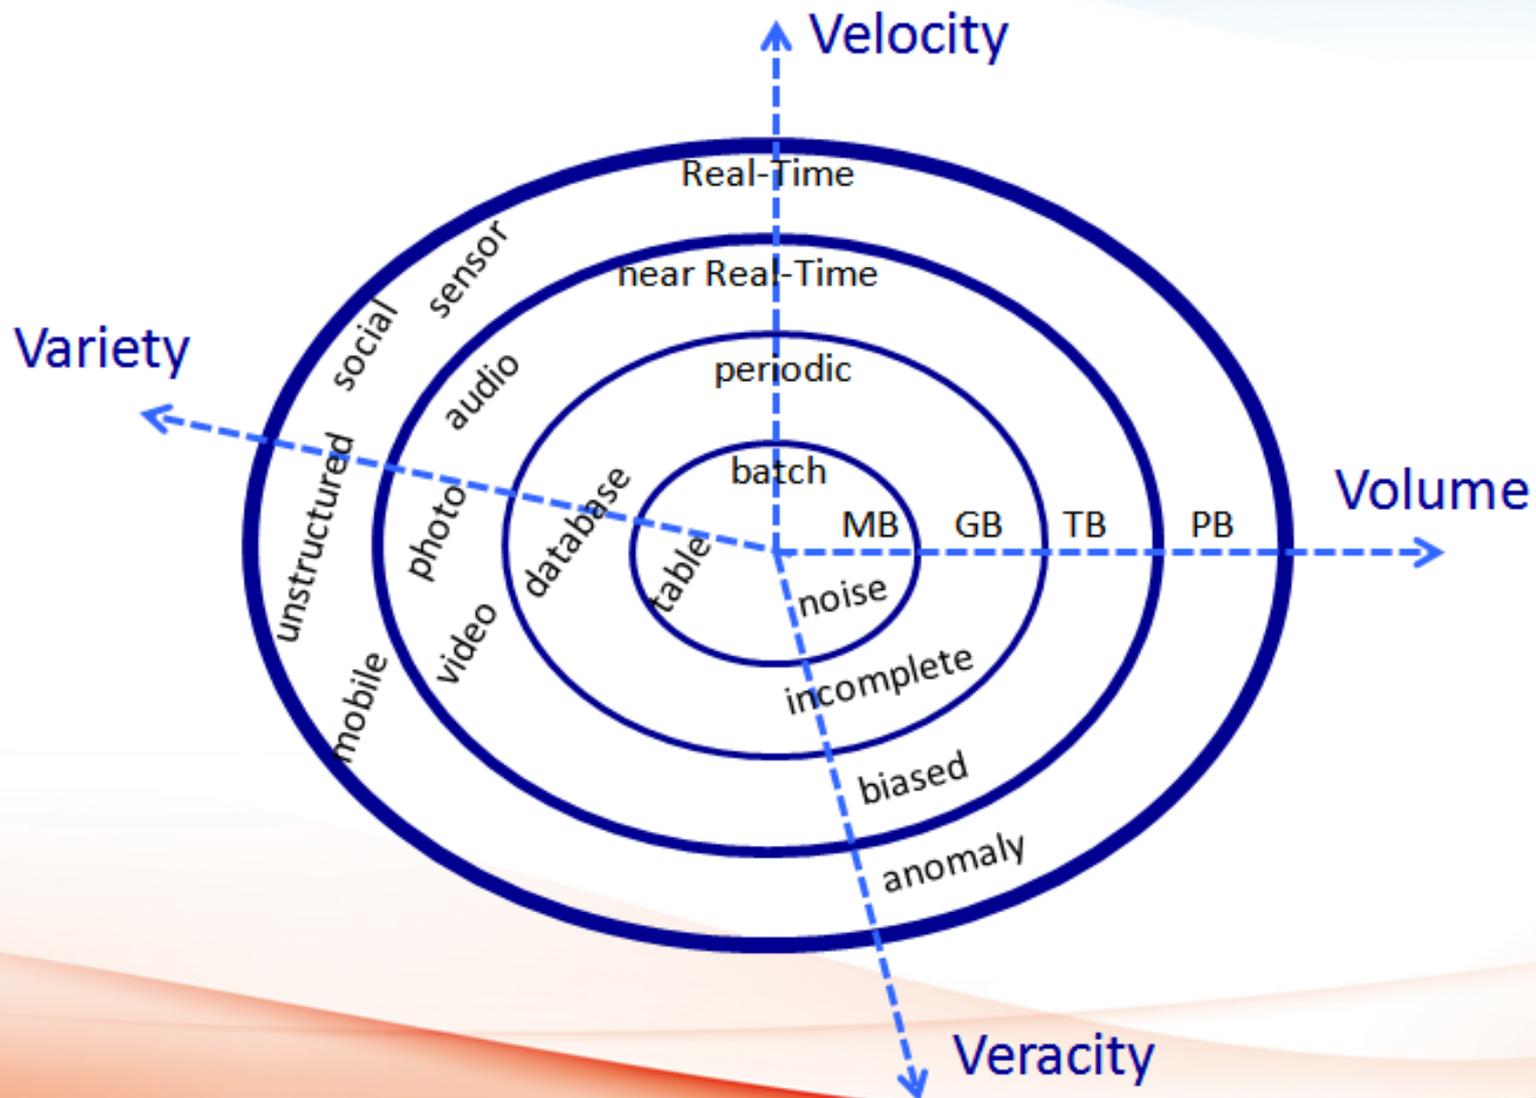


- What is Big Data?
- Big Data Technologies
  - Storage
  - Frameworks
  - Libraries
- Big Picture from IEEE Big Data Conference
  - Trends
  - Take away Lessons
- **Discussion**
  - **Future Directions for Fermilab**

# + What is Big Data?

- “**Big data** is a term describing the storage and analysis of large and or complex data sets using a series of techniques including, but not limited to: NoSQL, MapReduce and machine learning.”
  - 4Vs – **Volume, Velocity, Variety and Veracity**
  - big data is data which “exceed(s) the capacity or capability of current or conventional methods and systems.”
  - “Big data is the term increasingly used to describe the process of applying serious computing power—the latest in machine learning and artificial intelligence—to seriously massive and often highly complex sets of information.”
  - Big data opportunities emerge in organizations generating **a median of 300 terabytes of data a week**. The most common forms of data analyzed in this way are business transactions stored in relational databases, followed by documents, e-mail, sensor data, blogs, and social media.

# Big data Expands on 4 fronts



# SO WHAT IS A PETABYTE ANYWAY?

Source – www.mozy.com

## WHAT IS A PETABYTE?

TO UNDERSTAND A PETABYTE WE MUST FIRST UNDERSTAND A GIGABYTE.

**1** GIGABYTE = 7 MINUTES OF HD-TV VIDEO

**2** GIGABYTES = 20 YARDS OF BOOKS ON A SHELF

**4.7** GIGABYTES = SIZE OF A STANDARD DVD-R

THERE ARE A MILLION GIGABYTES IN A PETABYTE

*“Let me repeat that: we create as much information in two days now as we did from the dawn of man through 2003.” (That’s something like 5 Exabytes of Data). - Eric Schmidt – Google 8/10*

# A PETABYTE IS A LOT OF DATA

**1** PETABYTE = 20 MILLION FOUR-DRAWER FILING CABINETS FILLED WITH TEXT

**1** PETABYTE = 13.3 YEARS OF HD-TV VIDEO

**1.5** PETABYTES = SIZE OF THE 10 BILLION PHOTOS ON FACEBOOK

**15+** PETABYTES = INTERNET USER'S DATA BACKED UP ON MOZY.COM

**20** PETABYTES = THE AMOUNT OF DATA PROCESSED BY GOOGLE PER DAY

**20** PETABYTES = TOTAL HARD DRIVE SPACE MANUFACTURED IN 1995

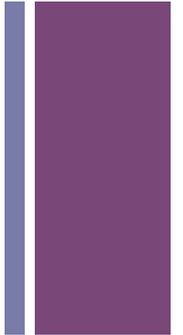
**50** PETABYTES = THE ENTIRE WRITTEN WORKS OF MANKIND, FROM THE BEGINNING OF RECORDED HISTORY, IN ALL LANGUAGES

Twitter:  
Over 7TB a Day in Tweets.

A ZETABYTE IS ONE MILLION PETABYTES!

Facebook:  
More that 750 Million Users.  
Average user creates 90 Pieces of content each month.  
More than 30B pieces of content shared each month.

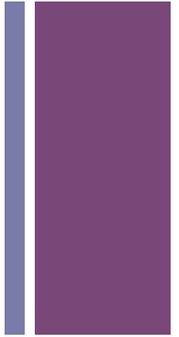
# + What comes with Big Data



- Where to store it
- How to store
- What to store
- How to retrieve/access
- How to process/query
- ...



# Big Data Technologies – Google

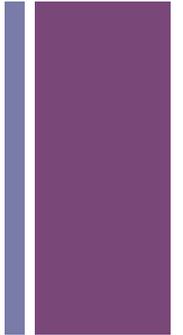


Here is a list of some of Google's proprietary software relating to big data:

- **GFS (Google File System)** – Google's scalable, fault-tolerant, distributed file system. Designed from scratch for use with data-intensive applications.
- **BigTable** – A distributed, high-performance database system built on top of GFS.
- **Spanner** – Google's scalable, multi-version, globally distributed, and synchronously replicated database. It is the first system to distribute data at global scale and support externally consistent distributed transactions.
- **MapReduce** – A framework for distributed processing of very large data sets.
- **Colossus** – A framework for real-time search.
- **Pregel** – A framework for analyzing large-scale graphs with billions of nodes.
- **Dremel** – Meant as a faster complement to Mapreduce, Dremel is a scalable, interactive, ad-hoc query system for large data sets. According to Google, it's capable of running aggregation queries over trillion-row tables in seconds and scales to thousands of CPUs.



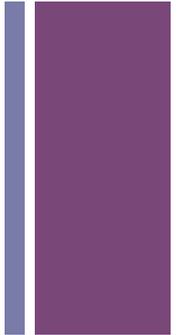
# Big Data Technologies – Storage



- **HDFS (Hadoop Distributed File System)** – Part of the open source Hadoop framework, HDFS is a distributed, scalable file system inspired by the Google File System. It runs on top of the file system of the underlying OSs and is designed to scale to petabytes of storage. The Hadoop project (you'll see several of the other components further down) has several high-profile contributors, the main one being Yahoo. Hadoop is used by Yahoo, AOL, eBay, Facebook, IBM, Meebo, Twitter and a large number of other companies and services.
- **CloudStore (KFS)** – An open source implementation of the Google File System from Kosmix. It can be used together with Hadoop and Hypertable. A well-known CloudStore user and contributor is Quantcast.
- **QFS** – based on the KFS code base, no fundamental difference, it is the newer version of KFS.
- **GlusterFS** – A free, scalable, distributed file system developed by Gluster



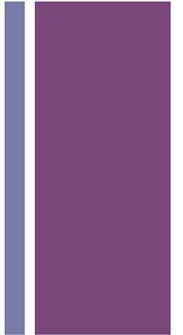
# Big Data Technologies – NoSQL Databases



- **HBase** – A distributed, fault-tolerant database modeled after Google’s BigTable. It’s part of the Apache Hadoop project, and runs on top of HDFS.
- **Hypertable** – An open source database inspired by Google’s BigTable. A notable Hypertable user is Baidu.
- **Cassandra** – A distributed key-value database originally developed by Facebook, released as open source, and now run under the Apache umbrella. Cassandra is used by Facebook, Digg, Reddit, Twitter and Rackspace, to name a few.
- **MongoDB** – An open source, scalable, high-performance, document-oriented database. It’s used by, among others, Foursquare, Bit.ly, Shutterfly, Etsy and Chartbeat.
- **Membase** – An open source, distributed, key-value database optimized for interactive web applications, developed by several team members from the famous Memcached project. Users include Zynga and Heroku. A month ago, the Membase project merged with CouchDB, creating a new project called Couchbase.



# Big Data Technologies – Querying and data analysis

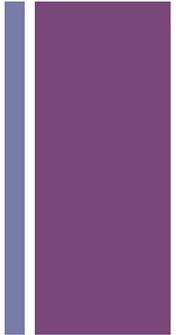


All that data is of no use without the ability to access, process and analyze it.

- **Hadoop MapReduce** – Open source version of Google’s MapReduce framework for distributed processing of large datasets.
- **Hive** – An open source data warehouse infrastructure with tools for querying and analyzing large datasets in Hadoop. Supports an SQL-like query language called Hive QL.
- **Pig** – A high-level language used for processing data with Hadoop. Funny aside: the language is sometimes referred to as Pig Latin.



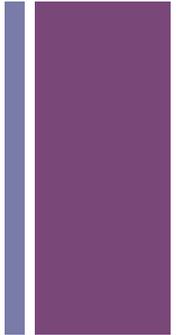
# Big Data Technologies – Data Processing Frameworks



- **S4** – A general-purpose, distributed, scalable platform for processing continuous streams of data. Developed by Yahoo and released as open source in 2010. It's apparently not quite ready for prime time yet, although Yahoo is using a version of it internally.
- **Esper** – An event-processing platform from EsperTech for handling continuous streams of incoming data.
- **StreamInsight** – Microsoft's entry in the EST/CEP field, included with SQL Server.



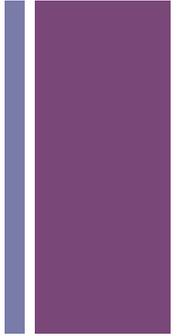
# Big Data Technologies – Data Processing Frameworks



Solution	Developer	Type	Description
Storm	Twitter	Streaming	Twitter's new streaming big-data analytics solution
S4	Yahoo!	Streaming	Distributed stream computing platform from Yahoo!
Hadoop	Apache	Batch	First open source implementation of the MapReduce paradigm
Spark	UC Berkeley AMPLab	Batch	Recent analytics platform that supports in-memory data sets and resiliency
Disco	Nokia	Batch	Nokia's distributed MapReduce framework
HPCC	LexisNexis	Batch	HPC cluster for big data



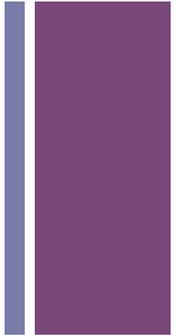
# Big Data Technologies – Libraries



- To allow for usage of efficient distributed machine learning and data mining algorithms, there are Hadoop based libraries available
  - Mahout
  - GraphBuilder – not as mature as Mahout



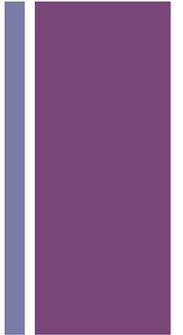
# Big Picture from IEEE Big Data Conference - Trends



- What People were talking about the most
- Keynote – The only relatively interesting one
- Title: **The Berkeley Data Analytics Stack: Present and Future**
- Message:
  - Gave yet another definition of Big Data – time money quality
  - Three aspects: algorithms, machine learning and crowd sourcing
  - SHARK and SPARQ
    - Interactive in-memory data processing
  - Open Source is your best friend → ?



# Big Picture from IEEE Big Data Conference - Trends



## ■ Sessions

### ■ Graph analysis

- Using Hadoop for graph and data mining, in-memory analysis

### ■ Image Similarity Search Using Hadoop

- TB image data

### ■ Programming Intel Phi using Hadoop

- Extended an existing multi-core Hadoop version and optimized for Phi architectures

### ■ Others

- Some high performance I/O techniques and experiences

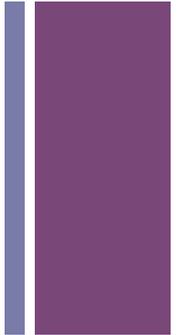
- Using Lustre FS for climate data

- Pnetcdf

- Collective I/O



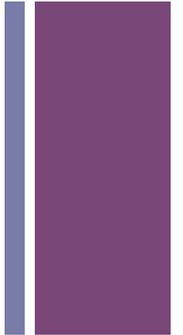
# Big Picture from IEEE Big Data Conference – Take away Lessons



- Big Data = Hadoop → Not a misconception
  - Primarily use for Storage
    - MapReduce for batch processing
    - SPARQL and others for Query Processing
  - All/Most problems that were presented can be solved using Hadoop and its variants
  - Most of the papers presented were either using Hadoop Or a modified version of Hadoop or some technology built on top of Hadoop e.g. graph mining, image similarity search, MapReduce for Phi.



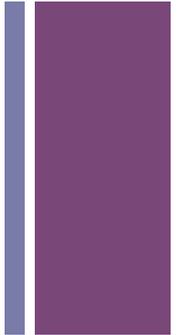
# Discussion – What is the Big Data Problem at Fermilab?



- Welcome to the "digital universe" — a measure of all the digital data created, replicated, and consumed in a single year. It's also a projection of the size of that universe to the end of the decade. The digital universe is made up of images and videos on mobile phones uploaded to YouTube, digital movies populating the pixels of our high-definition TVs, banking data swiped in an ATM, security footage at airports and major events such as the Olympic Games, **subatomic collisions recorded by the Large Hadron Collider at CERN**, transponders recording highway tolls, voice calls zipping through digital phone lines, and texting as a widespread means of communications.

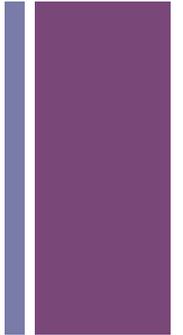


# Discussion – What is the Big Data Problem at Fermilab?



- CERN is building out a new data center and re-writing its tool chain to keep up with the big data demands of its research into the nature of matter and anti-matter.
  - Using PuppetsLab OpenStack for IaaS (local customized clouds?)
- [Link to Presentation - Attacking Big Data at CERN](#)

# + Summary



- What's out there
  - Storage:
    - File Systems
    - Database
  - Processing Frameworks:
    - Interactive/streaming/real-time
    - Batch
  - Libraries with distributed algorithms
    - Machine Learning
    - Data mining
- What we need