# Data Management for Climate Modeling

## The Present and the Future

V. Balaji, Amy Langenhorst, Aleksey Yakovlev

Fall Creek Workshop, Memphis, 2010

# Climate Modeling and Its Data (1/4)
## (the GFDL present)

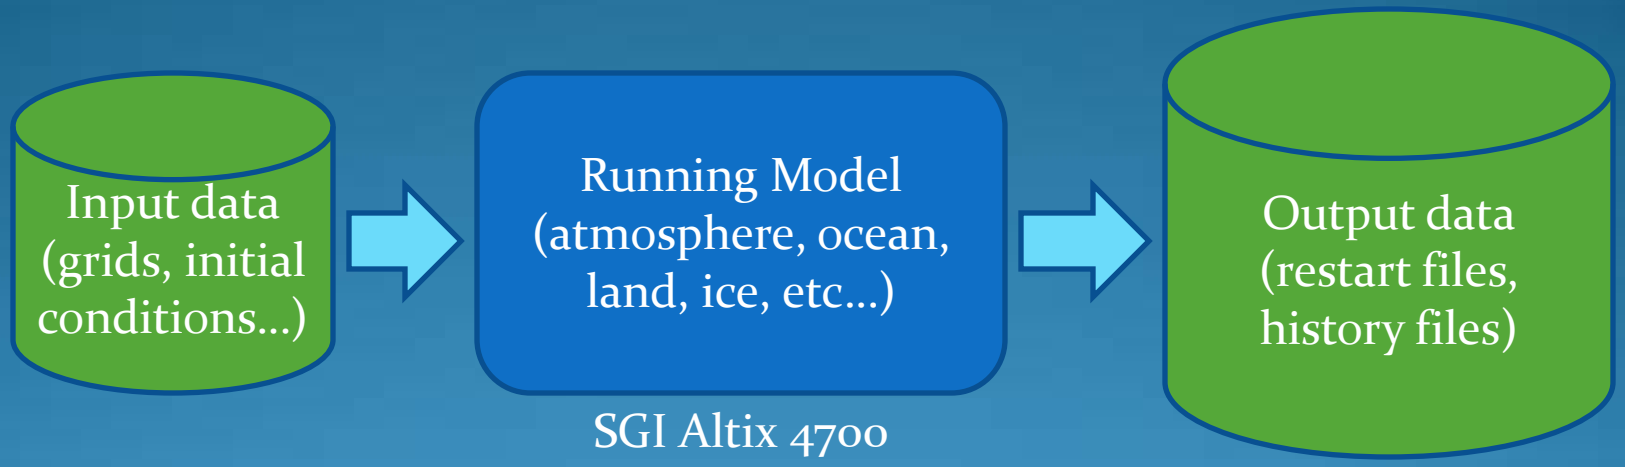


**Typical input data:**

- Number of files ~ **200**
- Total size ~ **10** GB

**Typical output data (per model year):**

- Number of files ~ **50**
- Total size ~ **20** GB
- Typical atmospheric grid – **144 x 90 x 2**
- Typical ocean grid – **360 x 200 x 50**

# Climate Modeling and Its Data (2/4)
## (the GFDL present)

**Typical post-processed data:**

- Number of files – **thousands**
- Total size ~ roughly **twice** raw output data

# Climate Modeling and Its Data (3/4)
## (the GFDL present)

**Tape Archive:**

- Total size now ~ **20000** TB
- Disk buffers ~ **200** TB

**Disk Storage ("ptmp"):**

- Total size ~ **650** TB

# Climate Modeling and Its Data (4/4)
## (the GFDL present)

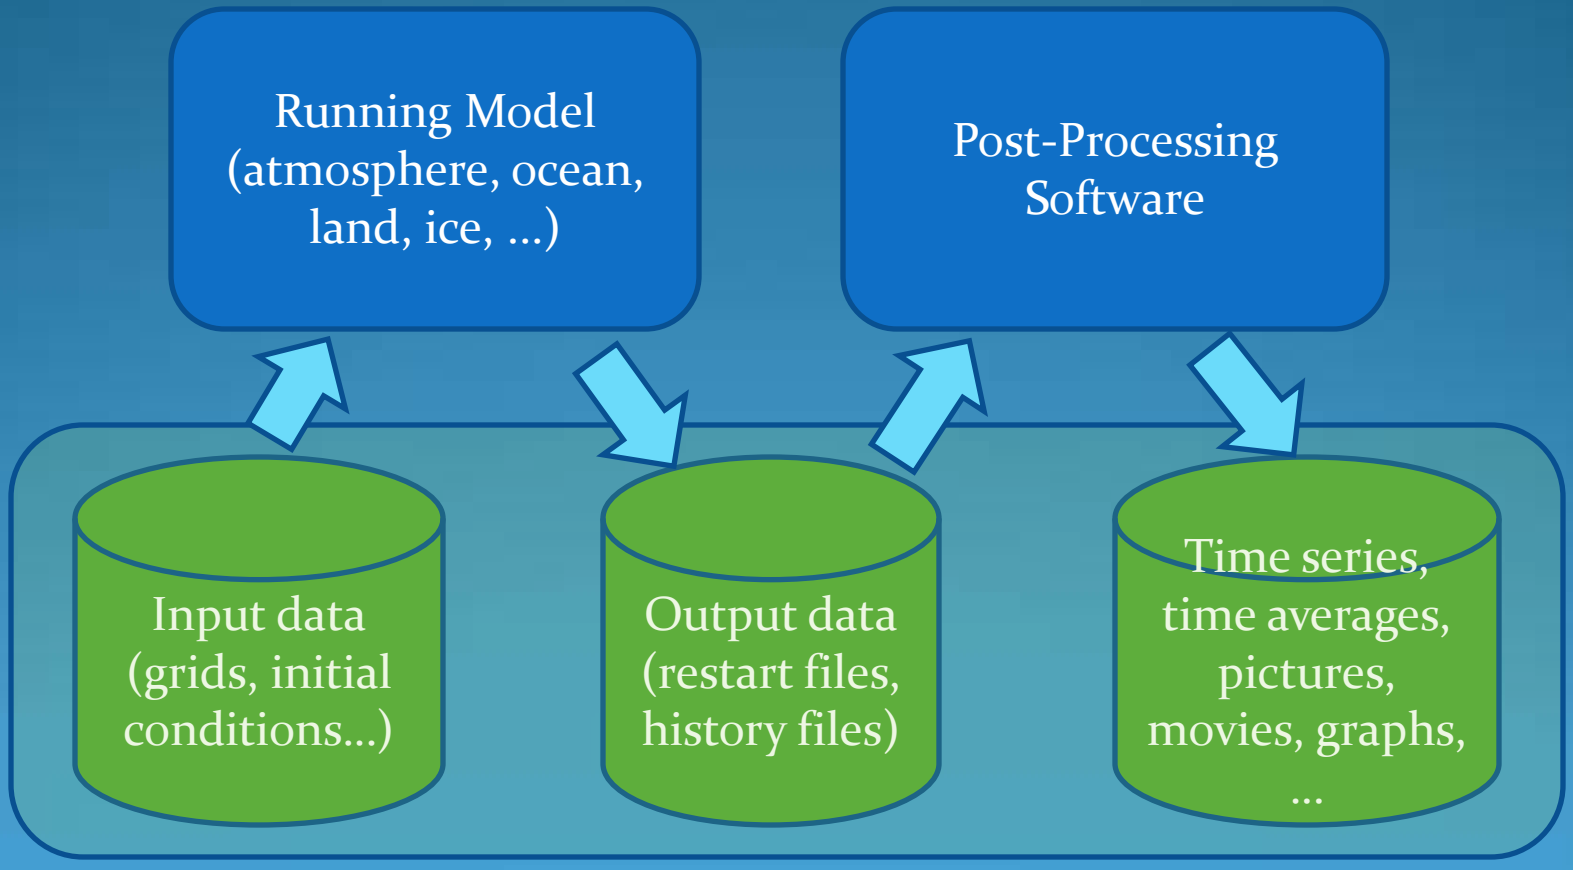


Everything starts and ends within the archive…

# Climate Modeling and Its Workflow Control (1/3) (the GFDL present)

Flexible Modeling System (**FMS**) is a software framework for supporting the efficient development, construction, execution, and scientific interpretation of atmospheric, oceanic, and climate system models

FMS Runtime Environment (**FRE**) is a toolset for managing experiments from start to finish. This includes tasks such as acquiring source code, compilation, launching jobs to run models, and post-processing the output

Hierarchical Storage Management (**HSM**) is a toolset for data management which allows FRE to work uniformly and efficiently with the data archive

# Climate Modeling and Its Workflow Control (2/3)
## (the GFDL present)

The HSM consists of two tools – "hsmget" and "hsmput":

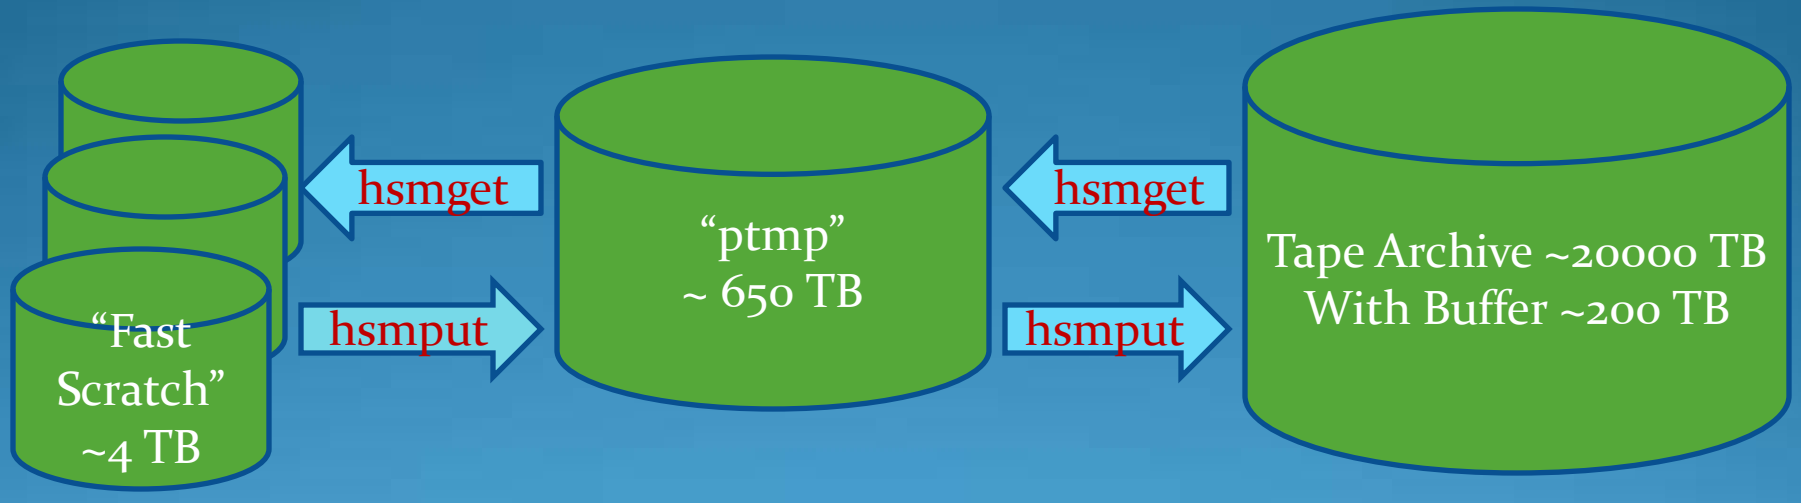


Both "hsmget" and "hsmput" call the GNU tool "make", which analyzes files timestamps and then makes decisions what data transfers are needed

# Climate Modeling and Its Workflow Control (3/3) (the GFDL present)

The intermediate data storage ("ptmp") is similar to cache:

- ✓ Cache hits/cache misses – yes
- ✓ Temporal locality – yes
- ✓ Spatial locality – no …
- ✓ Write policy - yes
- ✓ Replacement policy - ???

The HSM works with the intermediate "ptmp" data storage as with a cache, but not all the cache functionality can be efficiently implemented on a disk array with that size …

# Climate Modeling on Gaea (1/3)
## (the GFDL future)

Total number of cores ~ **30000**
Data storage: FS ~ **950** TB, LTFS ~ **3500** TB
Archive – **no**

# Climate Modeling on Gaea (2/3)
## (the GFDL future)

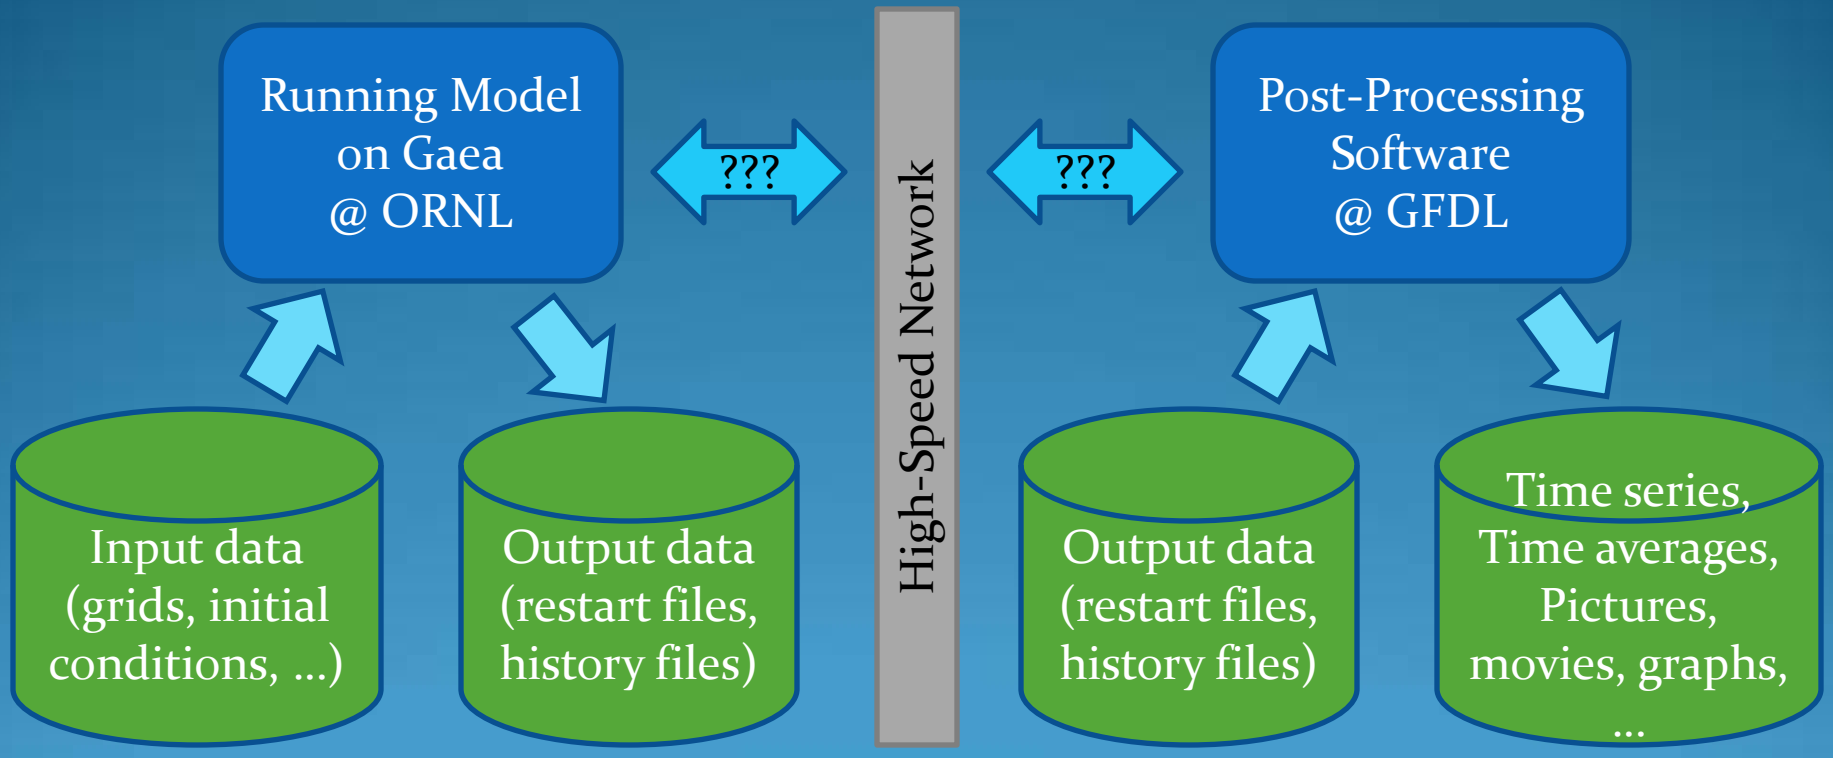


More complex architecture, but it will allow us to simulate better ….
typical atmospheric grid **720 x 360**, oceanic one – **1440 x 720**

# Climate Modeling on Gaea (3/3)
## (the GFDL future)

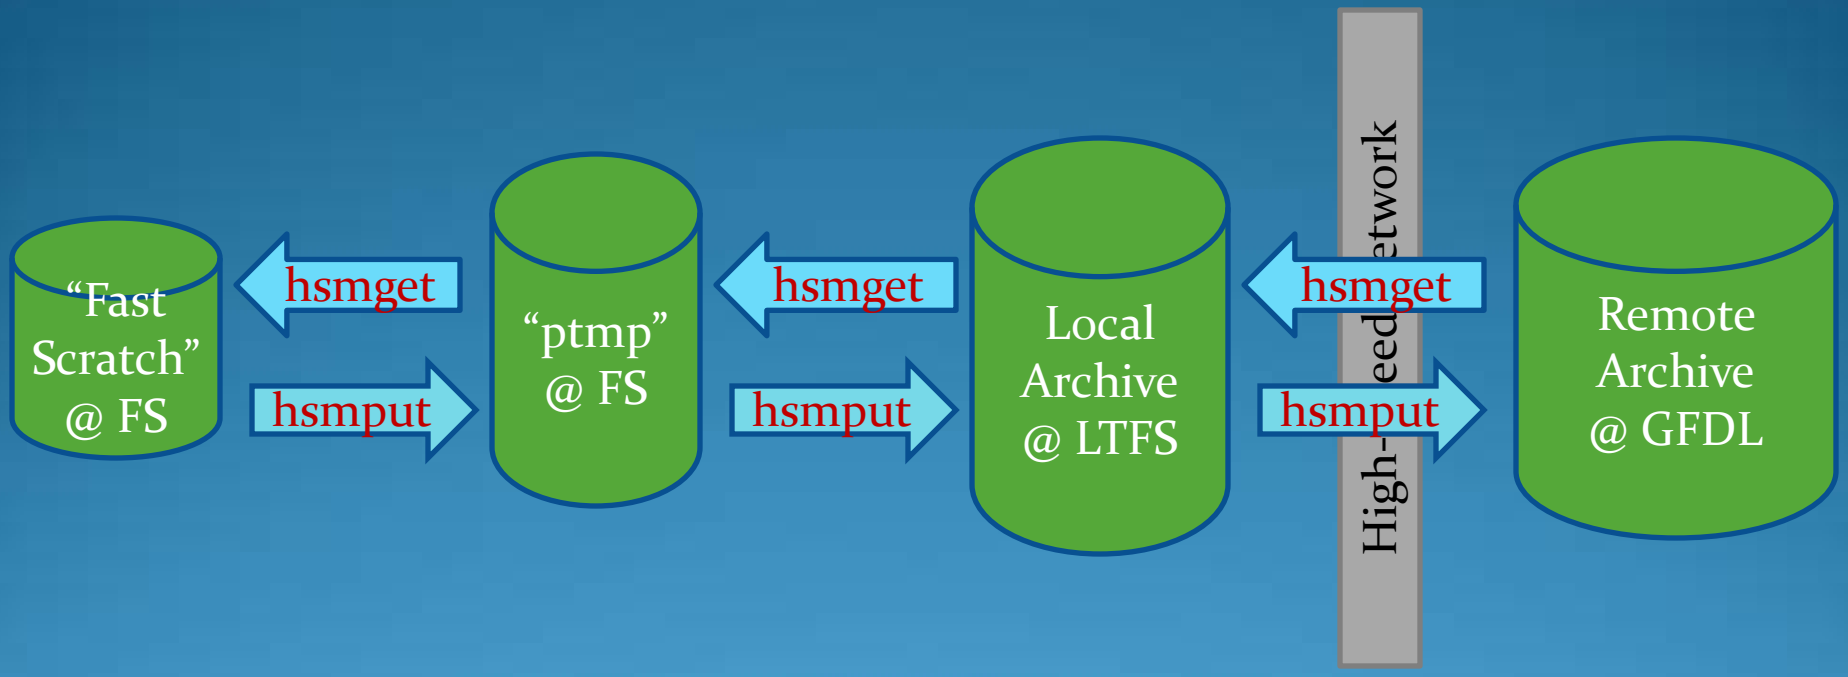


HSM gets one more caching level

HSM should work via a network => no easy way to compare timestamps

# “Common Sense” Data Management (1/3)
## (Multiple Copies of the Same Data)

Independent updates => manual merging
You left your laptop at home => access isn’t guaranteed
Your disc crashed => data is lost
Your laptop got stolen => your data isn’t private anymore

# “Common Sense” Data Management (2/3)
## (Single Copy of the Data)

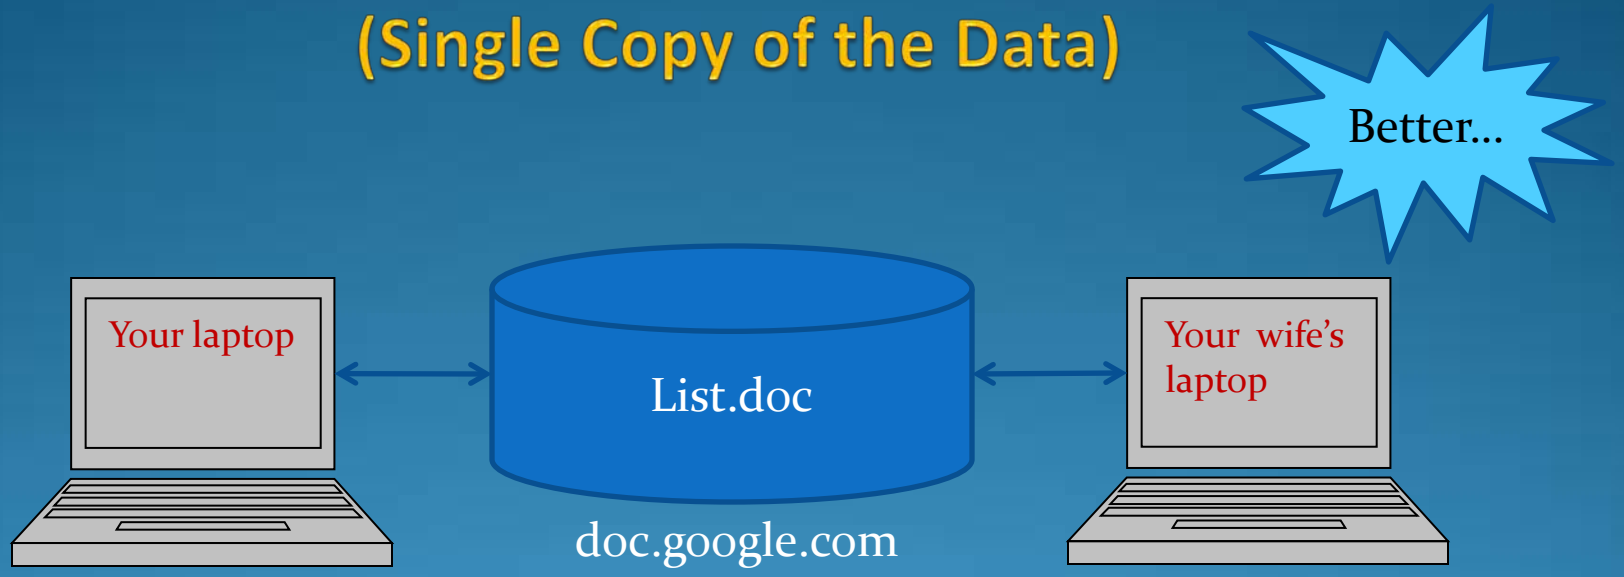


Single copy => no merging needed
You are using (almost) any computer => data is accessible
Your disk crashed => data isn’t lost
All your laptops got stolen => data is still available

(Editing might be slow and unusual, and what about privacy?)

# "Common Sense" Data Management (3/3)
## (All Versions of Your Data Under a Single Roof)

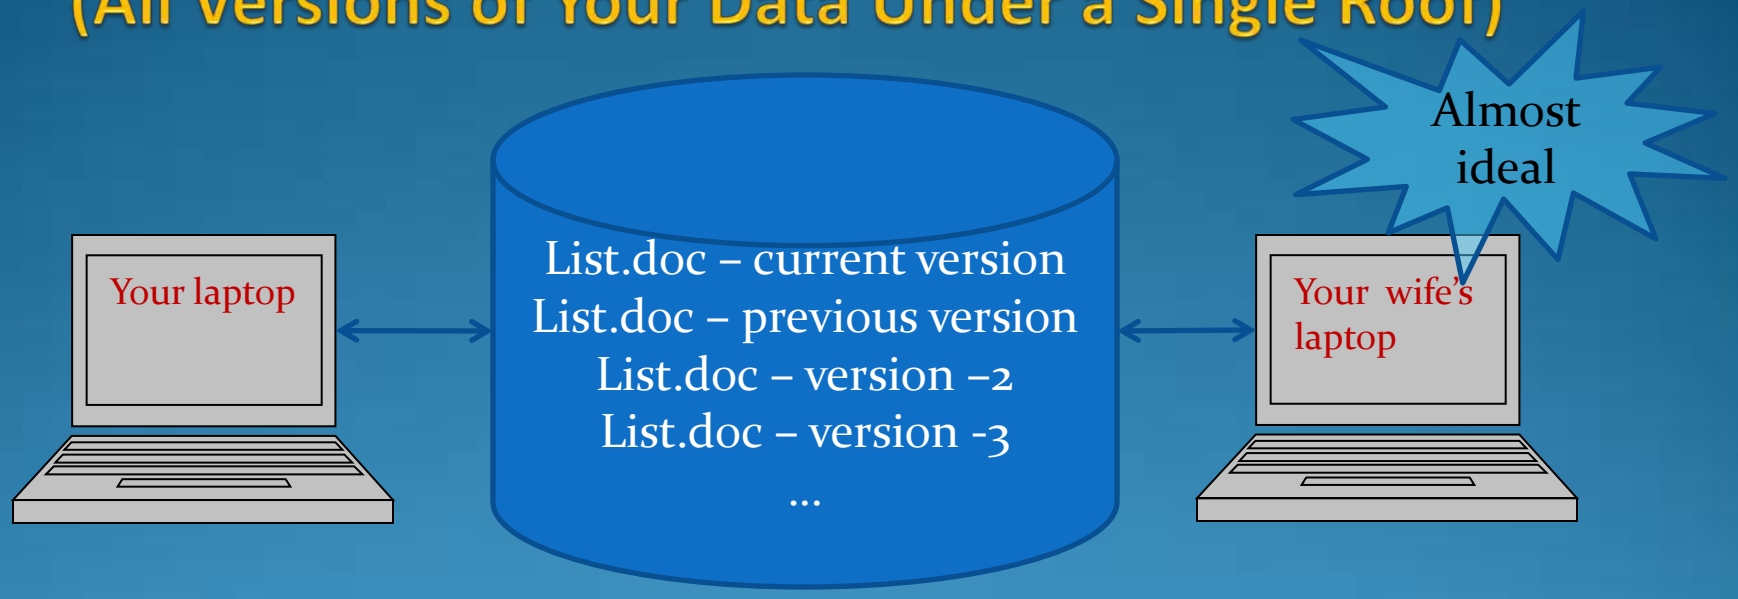


Version Control System (CVS, Subversion, etc…)

Controllable independent updates => (automatic) merging is possible
Comments, tags, update history => generalized view about data

(But what if data isn't textual?)

# What Do Our Users Want? (1/3)

- Do they want to manipulate their data themselves? Probably not
- Do they care about physical location of the data? No
- Do they care about data accessibility? Yes!
- Do they want to be sure that data won't be lost? Absolutely!
- Do they care about privacy of their data? To some extent…
- Do they want a single namespace for filenames? Yes
- Do they want versioned data storage? Yes
- Do they want metadata about their data? Oh, yes… it's a dream…

# What Do Our Users Want? (2/3)

**Climate modeling data specifics:**

- Write - once, read – many times, update – never
- Most of the data are multidimensional arrays of real numbers
- Size : files less than 1 gigabyte are considered small
- Data privacy isn't the main concern
- Need to access data is sporadic

# What Do Our Users Want? (3/3)

*Software technology to enable more intelligent data management in distributed computing environment*

**Our choices:**

- Distributed filesystems? (Hadoop etc.)
- Grid Computing? (Globus Toolkit etc.)
- others?

**Show stoppers:**

- Inter-organizational activities are always difficult =>
  - Security issues
  - Performance issues
  - Support issues

# References:

http://gfdl.noaa.gov/fms - GFDL Flexible Modeling System (FMS)
http://data1.gfdl.noaa.gov – GFDL Data Portal

Thank you!

Questions?