Thoughts on Computing Upgrade Activities

Origins

These thoughts primarily came from

- The White Papers contributed to the CWP activity by Maria Girone, Oliver Gutsche, Liz Sexton-Kennedy, and myself
- Some ideas that have been batted about

Includes

- General ideas on architectures
- Technology demonstrators we should be doing

Challenges

The challenge of the HL-LHC is about data

- Data Complexity and date rate both increase

However from a computing resource perspective some elements are more challenging than others

- Storage is much closer than the expected technology increase than processing
  - 6-8 times increase needed in storage but potentially 20 times increase in processing
    - This is because the processing times for complex events grows faster than their size, which tends to be linear with pile-up

Impact of this is we need to work harder at how we process than simply how we store

Model

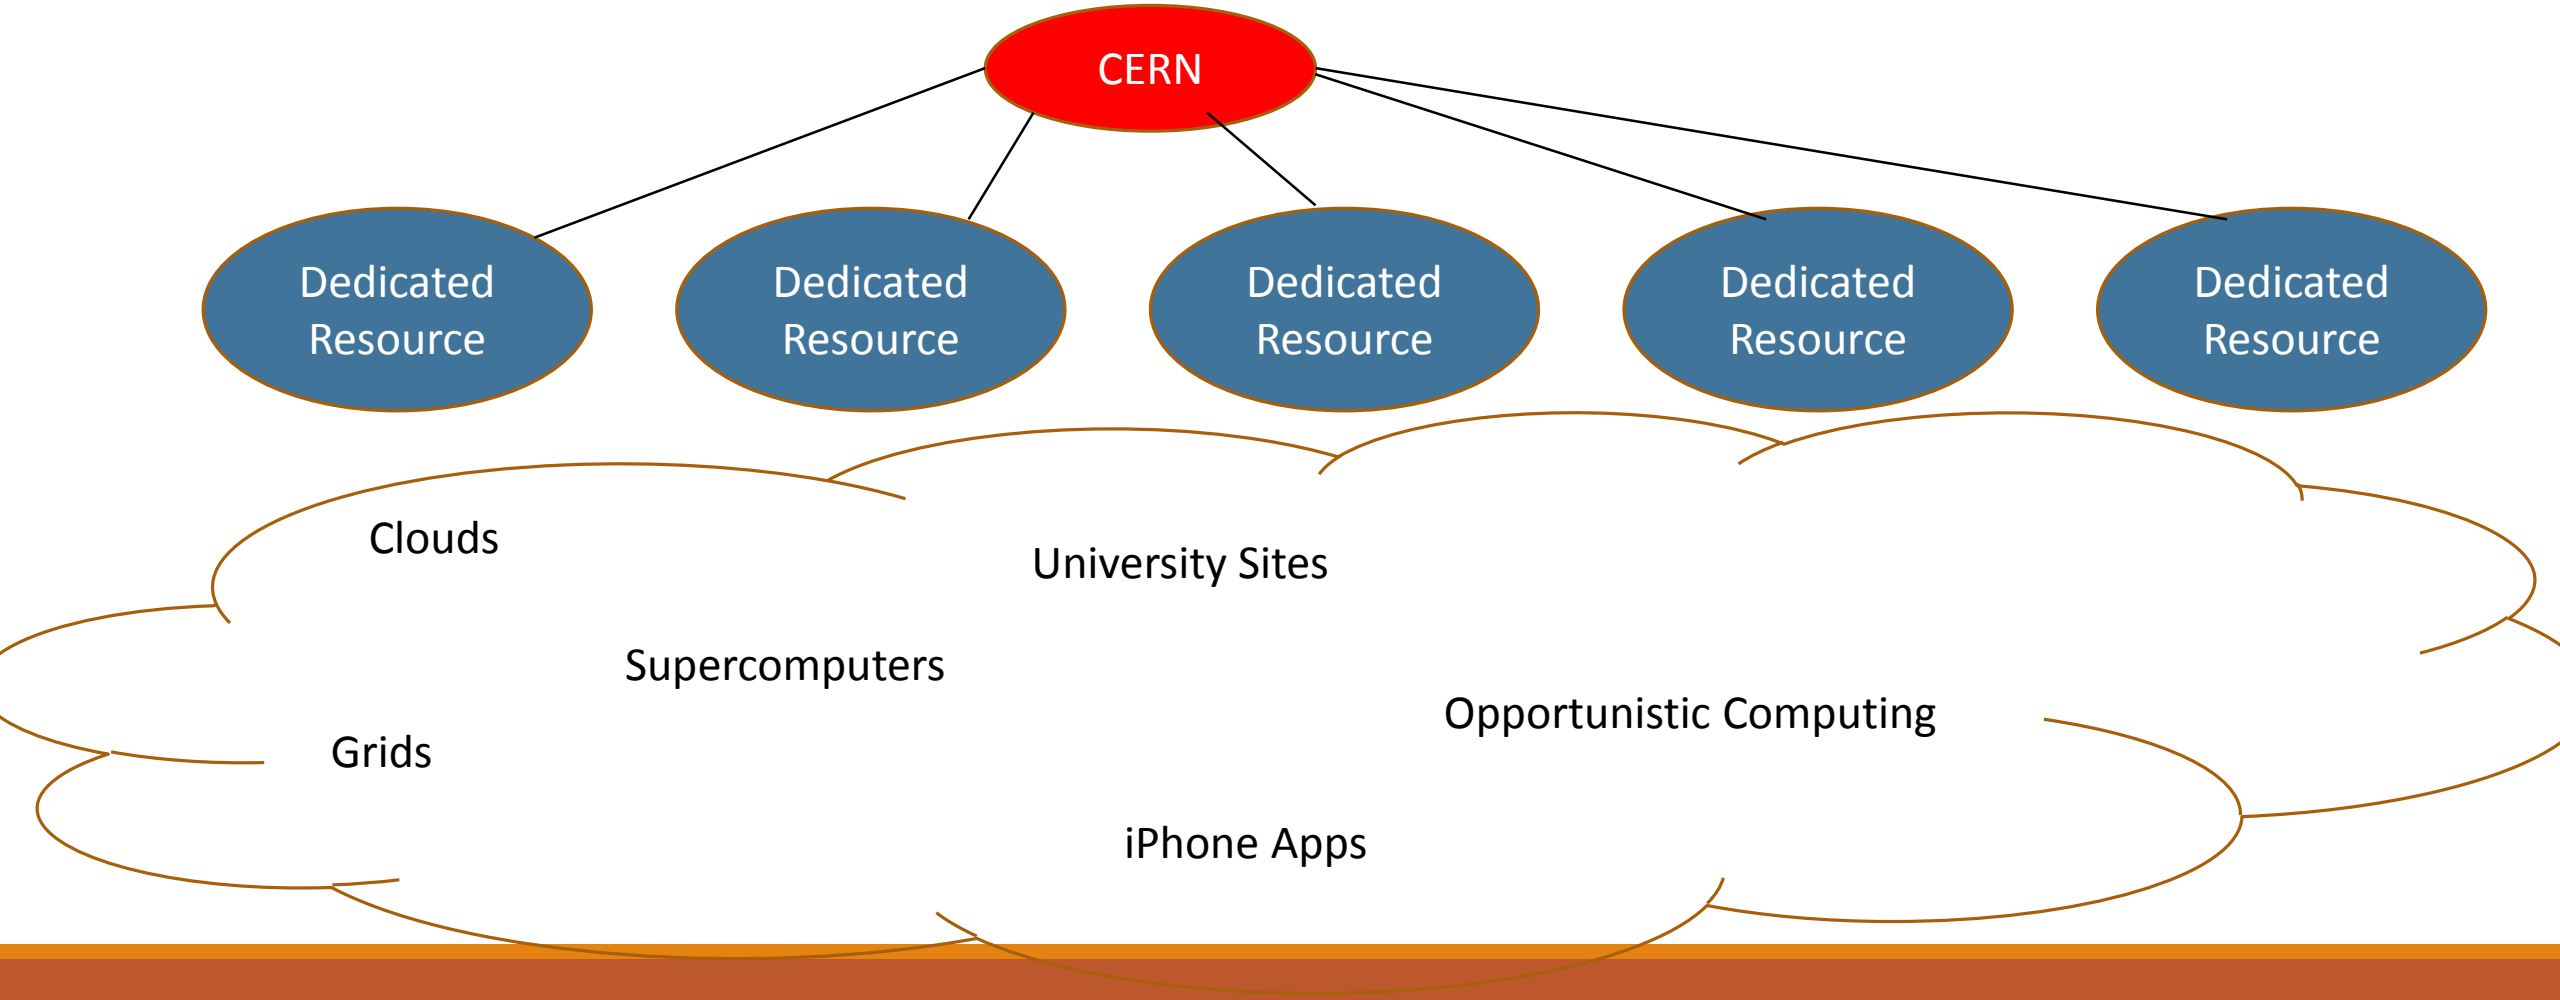
CERN
Dedicated Resource
Dedicated Resource
Dedicated Resource
Dedicated Resource
Dedicated Resource
Clouds
University Sites
Supercomputers
Opportunistic Computing
Grids
iPhone Apps

Data Replication

The whole system relies on the ability to move data
We have relied on GridFTP for data movement
- TCP requires many streams to get reasonable transfer rates
  - Pushes toward a program with parallel capabilities like GridFTP

Do we want to look at alternatives?
- UDP with checksums?
- Socket based connections?
- Commercial offerings like Aspera?

Potential Technology Demonstrator would be to look at an alternative to GridFTP
- Firewall friendly, deployable in many environments, capable of filling a 100Gb/s link over a reasonable latency connection

Resources: Storing Data

It is not obvious to me that on the time scale of HL-LHC that we should imagine data ever comes off tape except for disaster recovery

- ORACLE is getting out of the tape business
- We probably still want a tape copy, but we probably don’t care that all large sites have tape
- There will be fewer manufacturers of tape drives
- We will still have tiered storage
  - Demonstrator use erasure code written object storage as a cheap intermediate disk layer

Distributing Data

Since much of this model would rely on making effective use of a variety of resources away from where the data is permanently stored and served from. We think we need demonstrators on data management

- Named Data Networks (NDN) combined with Software Defined Networks (SDN) to ensure that a group of files can arrive at a remote site for processing
  - Integrate meta data and data relationships with network layer
  - This is a potential project with the next phase of Openlab
- Improved data federation and smarter caching
  - Demonstrate use of sites with no dedicated storage and only caching
  - Evolve FAX and AAA in common
  - Development of smarter caches
    - Similar elements to NDN

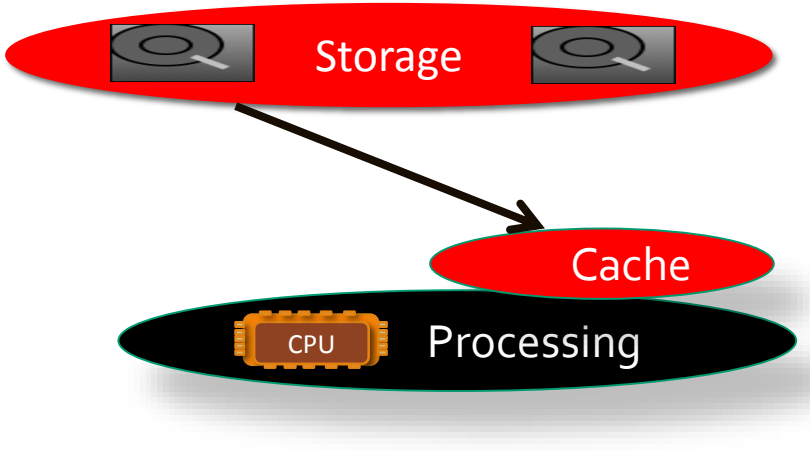

Resources: Processing

A lot of the resources we rely on to close the gap will not be dedicated processing we bought

- Will be cloud computing that was donated or we payed for
  - Engagements with commercial cloud providers have gone well
- Will be HPC shares
- Will be opportunistic

There is no reason to think these resources will or should come with a flat profile

- Technology demonstrator we would like to do with Condor
  - 10M processing jobs executed
  - Condor executing “work” and ”job” scheduling is under the hood

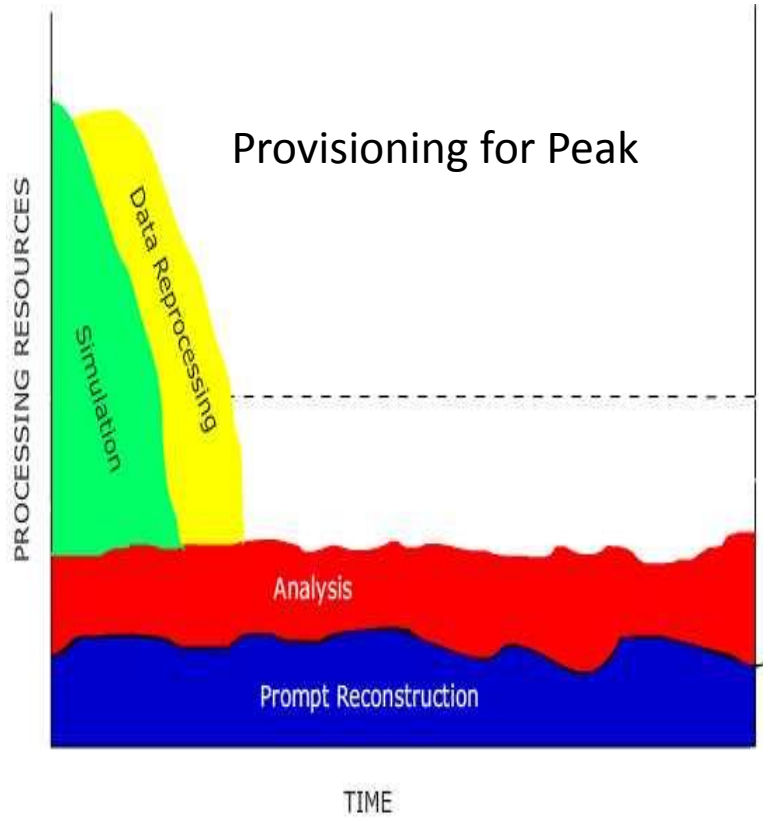

Physics Data Reduction

Dealing with the data and the wide area gets a lot simpler if there is less data

- Using data analytics tools to support data selection and reduction
- Goal is to use map/reduce type selections to take PBs and make them TBs for more detailed offline analysis
- On-going demonstrator with Intel through openlab

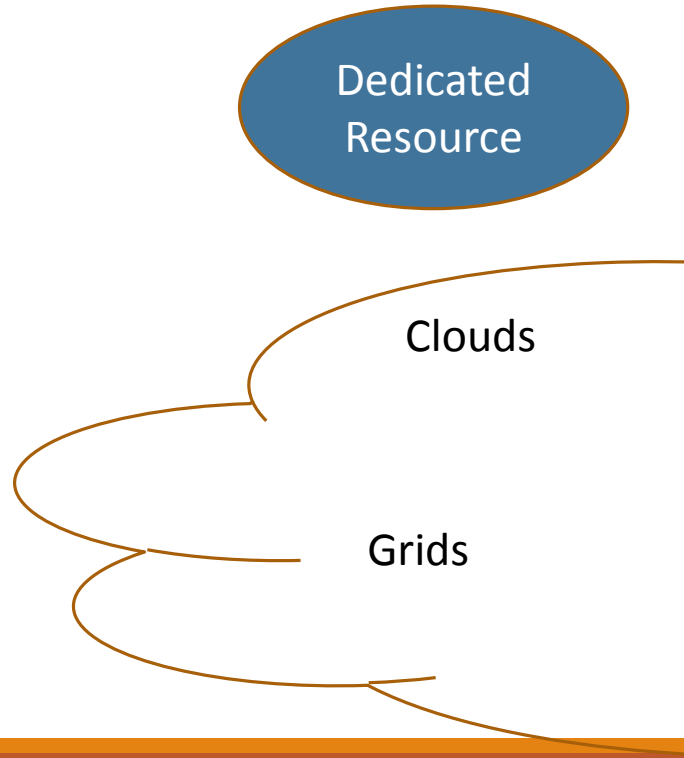

Data Transformation

The simple cut based data reduction might evolve into more advanced data refresh

- Accessing a data quantity could trigger recalculating it
- See data reconstruction not necessarily as a big event loop, but as a series of discreet data transformations
  - Look at new programming languages and paradigms

New Project with the EU called DEEP-EST using HPC resources would be an interesting place to test some of these ideas

Demonstrators

Data Transfer

Use of Object Store for the layer between disk and tape

Named Data Networks

Advanced Caching

Scalable Workflow Scheduling

Data Reduction

Data Transformation

Outlook

I think we have some techniques to look at and the basics of a computing model general enough to support almost anything

I don’t think we know how much any of the technologies we will investigate will actually go toward closing the resource gap