CMS Cloud Plans

Cloud Work

- CMS recently performed an EC2 study
  - We successfully configured and ran a Monte Carlo production on the cloud
    - Configured services and ran simulation workflows on the cloud
    - Generally successful
    - performed a comparison of the relative cost of EC2 to dedicated computing
      - Factoring in systems, power, cooling, admin costs we determined that EC2 was about 8 times more expensive than locally provided computing assuming it is used regularly and efficiently

Commercial Clouds

- While it's useful to have the flexibility to utilize commercial clouds if the market changes, currently we see it as only peak
  - If you need short term computing and can't get it anywhere else, this is an attractive option
  - Hard to imagine that it will ever be cost competitive for computing used with high efficiency

Cloud Scheduling

- In CMS we think the cloud scheduling is an interesting technology
  - We believe the hard problems to make it work are not in scheduling but in all the other services
    - Fortunately, most of the work is needed to improve regularly scheduling also
      - Glide-In WMS Pilots
      - Whole Node Scheduling
      - Remote access to data
      - Instantiated services

Whole Node Scheduling

- CMS is interested in a resource allocation mechanism that allows executing a process in user space on a whole multi/many-core host
  - Clear benefits in term of memory consumption
  - Still to be understood the implications for the local I/O
  - We expect to have more answers by the work of the Whole Node Task Force
- We think whole node matches Cloud Scheduling well
  - Persistent appliances scheduled for longer periods

Data Access Plans

- Attempting to deploy a regional infrastructure to provide limited access to data remotely
  - xrootd based with a European and a US region (so far), and eventually a global redirector on top
  - Uses existing storage systems but with xrootd on top
  - Number of sites participating is limited as sources and clients is limited but growing
  - 3 Initial Use Cases are targeted
    - 1.) Backup channel for local storage
      - Failure to open a file could fall back to remote access
        - Data popularity provides statistics on failures to open
    - 2.) Visualization
      - Idea is that any event should be available to visualize
    - 3.) Debugging
      - Augment the FileMover functionality by providing access to data files to applications
        - Potentially the most risk of abuse and throttles and monitoring are needed.

Remote Access

- Especially for opportunistic use of commercial clouds breaking data locality is helpful
  - A lot of improvements in the IO layer
    - Better reads of only the data needed
    - Ordering of reads
    - Caching