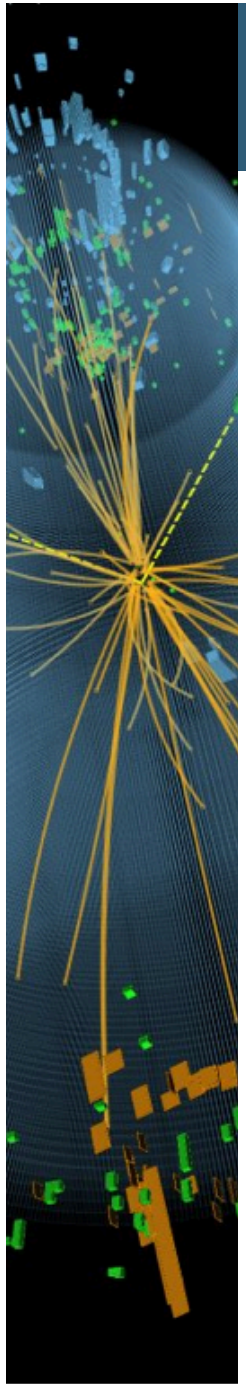

# Common Solutions for LHC Computing Problems

Ian Fisk, Maria Girone, and Alexei Klimentov

On behalf of the CMS, CERN IT-ES/VOS and ATLAS

# Introduction

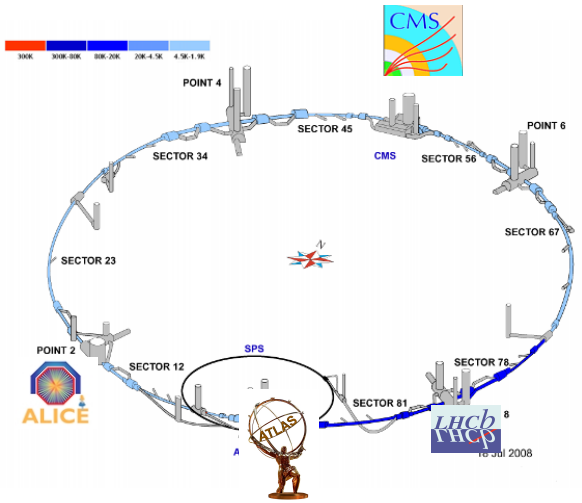


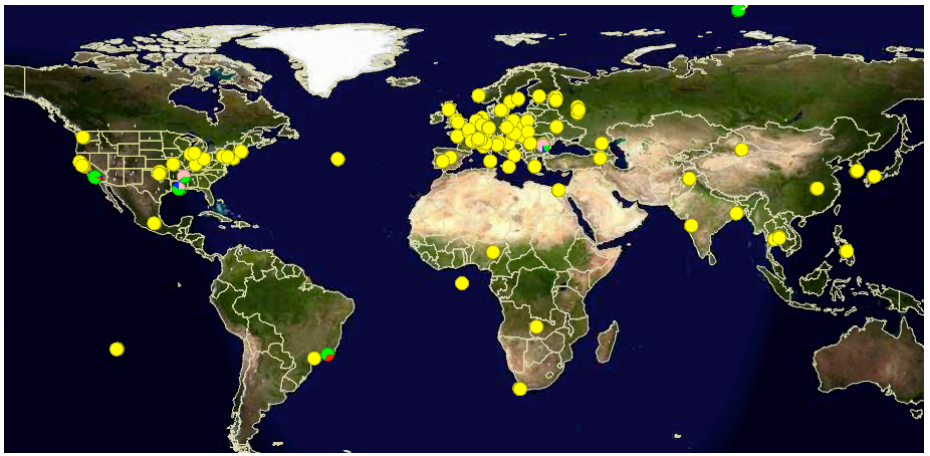

- Two of the LHC experiments are general purpose, and two are more specialized
  - From a computing perspective a lot of the workflows are similar and can be done with common services
- While the experiment collaborations are huge and highly distributed, effort available in development is limited and decreasing
  - Effort is focused on analysis and physics
    - Common solutions are a more efficient use of effort

Ian Fisk
FNAL/CD

# Anatomy of the Common Solution


Most common solutions can be diagrammed as the interface layer between common infrastructure elements and the truly experiment specific components

- One of the successes of the grid deployment has been the use of common grid interfaces and local site service interfaces
- The experiments have a environments and techniques that are unique
- In common solutions we target the box in between. A lot of effort is spent in these layers and there are big savings of effort in commonality
  - not necessarily implementation, but approach & architecture

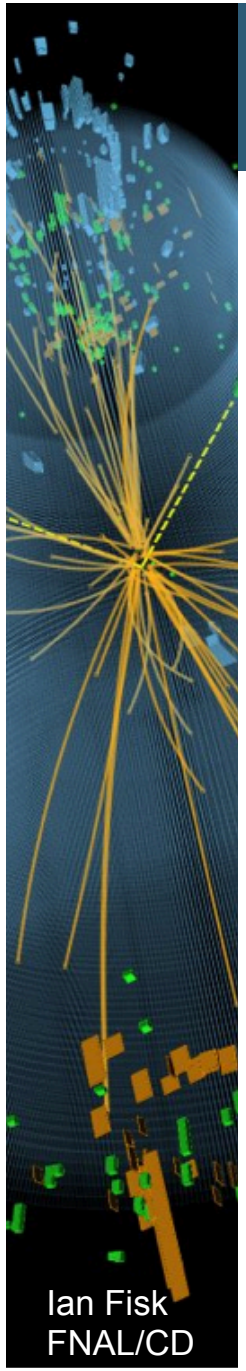


# The Group

- Experiments have a history of using common components through the grid interfaces
- In this project we rely on expertise from the experiments and IT-ES/VOS
  - The group is currently supported with substantial EGI-InSPIRE project effort
  - Careful balance of effort embedded in the experiments & on common solutions
  - Development of institutional expertise in experiment systems across experiment boundaries
  - People uniquely qualified to identify and implement common solutions
    - Matches well with the EGI-InSPIRE mandate of developing sustainable solutions

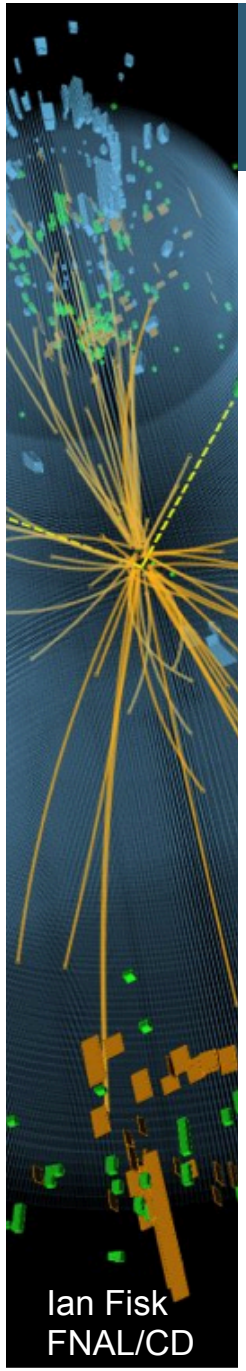


# Examples

- Data Management support
  - Covers development and integration of the experiment specific and shared grid middleware
- Monitoring and Experiment Dashboards
  - Allows experiments and sites to monitor and track their production and analysis activities across the grid
- The LCG Persistency Framework
  - handles the event and detector conditions data from the experiments
- Distributed Production and Analysis
  - design and development for experiment workload management and analysis components

# Example: Data Popularity

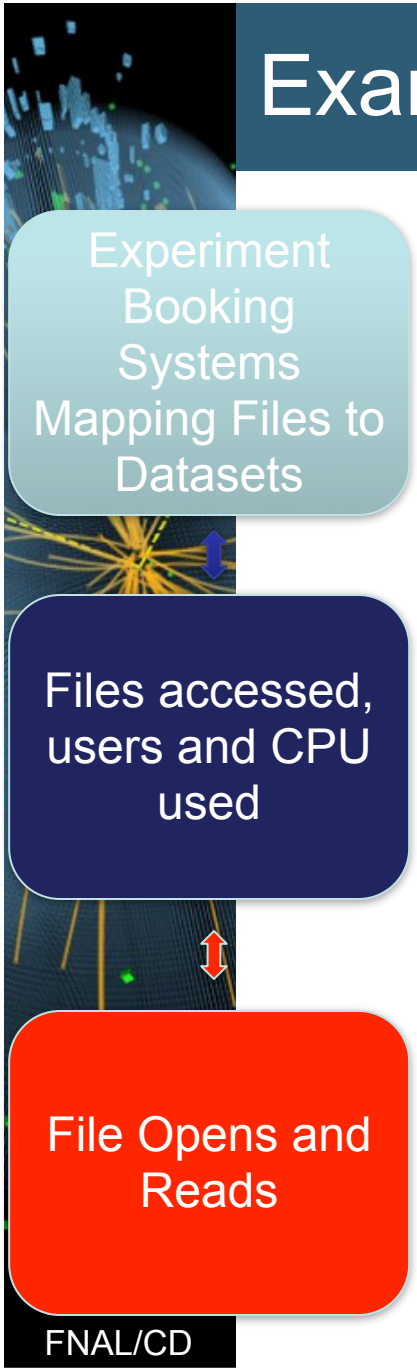


- The experiments have system that identify how a low level object like a file is mapped to a higher level logical object like a dataset
- All experiments open files
- Experiments want to know how the logical concepts like datasets are used, how much, and by whom
  - Good chance of a common solution

# Popularity Service

- Used by the experiments to assess the importance of computing processing work, and to decide when the number of replicas of a sample needs to be adjusted either up or down and replicate or clean-up

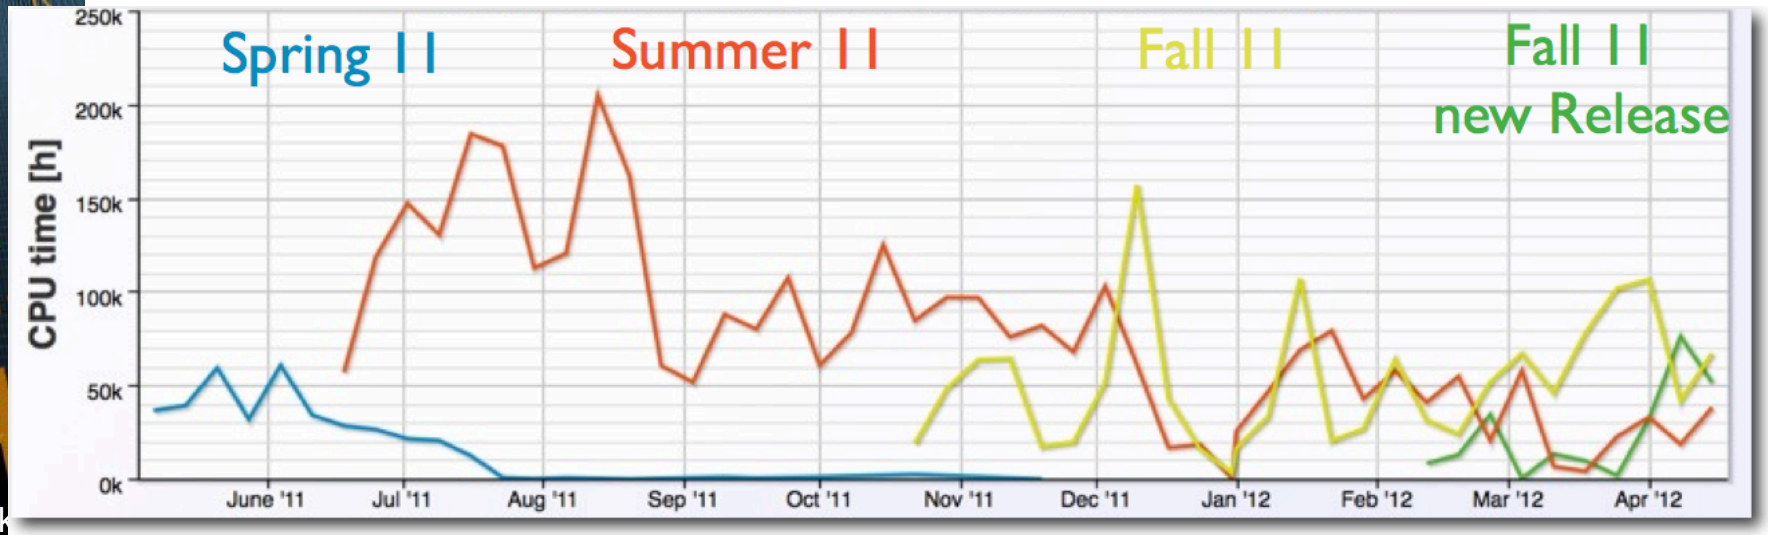


Time evolution of W+jet datasets

# New Activities – Analysis Workflow


- Up to now services have generally focused on monitoring activities
  - All of these are important and commonality saves effort
  - Not normally in the core workflows of the experiment
- Success with the self contained services has provided confidence moving into a core functionality
  - Looking at the Analysis Workflow

# Analysis Workflow Progress

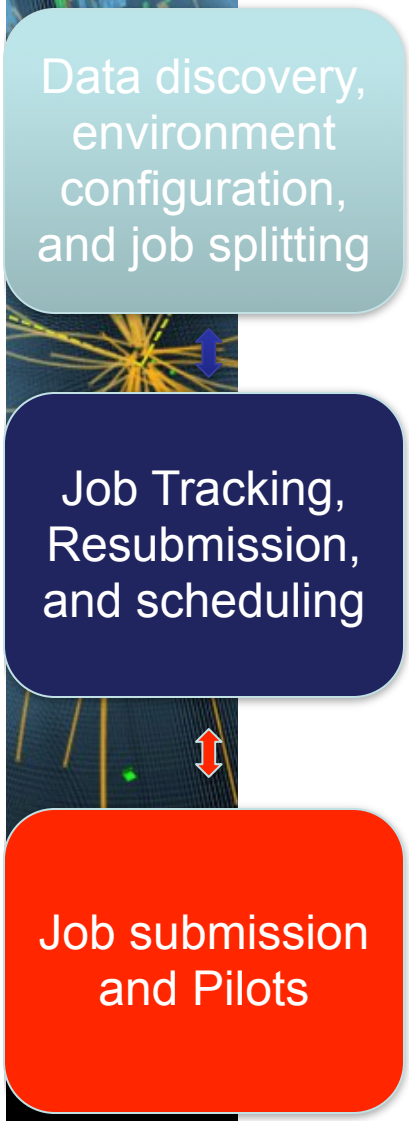


- Looking at ways to make the workflow engine common between the two experiments
  - Improving the sustainability of the central components that interface to low-level services
    - A thick layer that deals with tracking jobs after they are created (resource assignment, job tracking, resubmission)
  - Maintaining experiment specific interfaces
    - Job splitting, environment, and data discovery would continue to be experiment specific

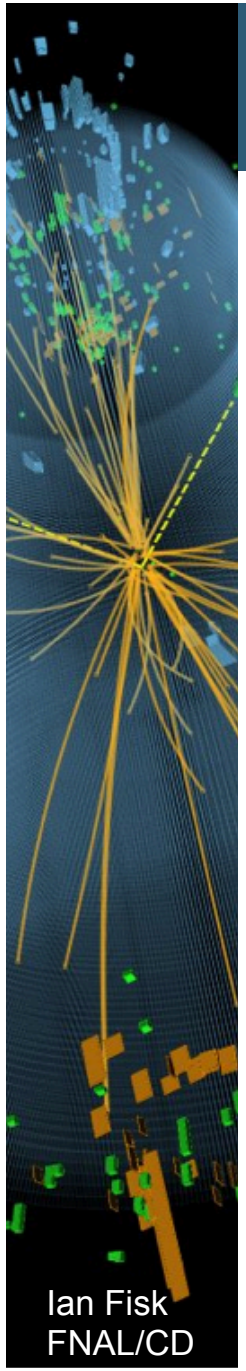


# Goal

- Take elements of both experiment systems, and combined with clean interfaces to experiment specific elements to develop a common solution
  - Workflow tracking from Panda
  - Pilot submission from Glide-in WMS
  - Job splitting and data discovery from the experiment elements
- Investigate scalability and functionality

# Proof of Concept Diagram

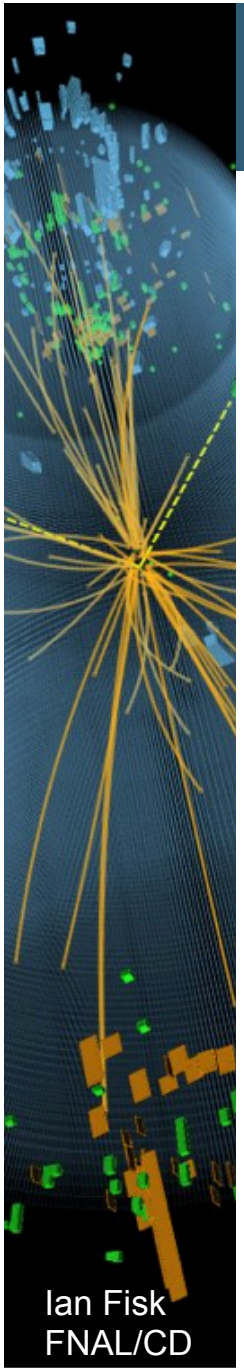

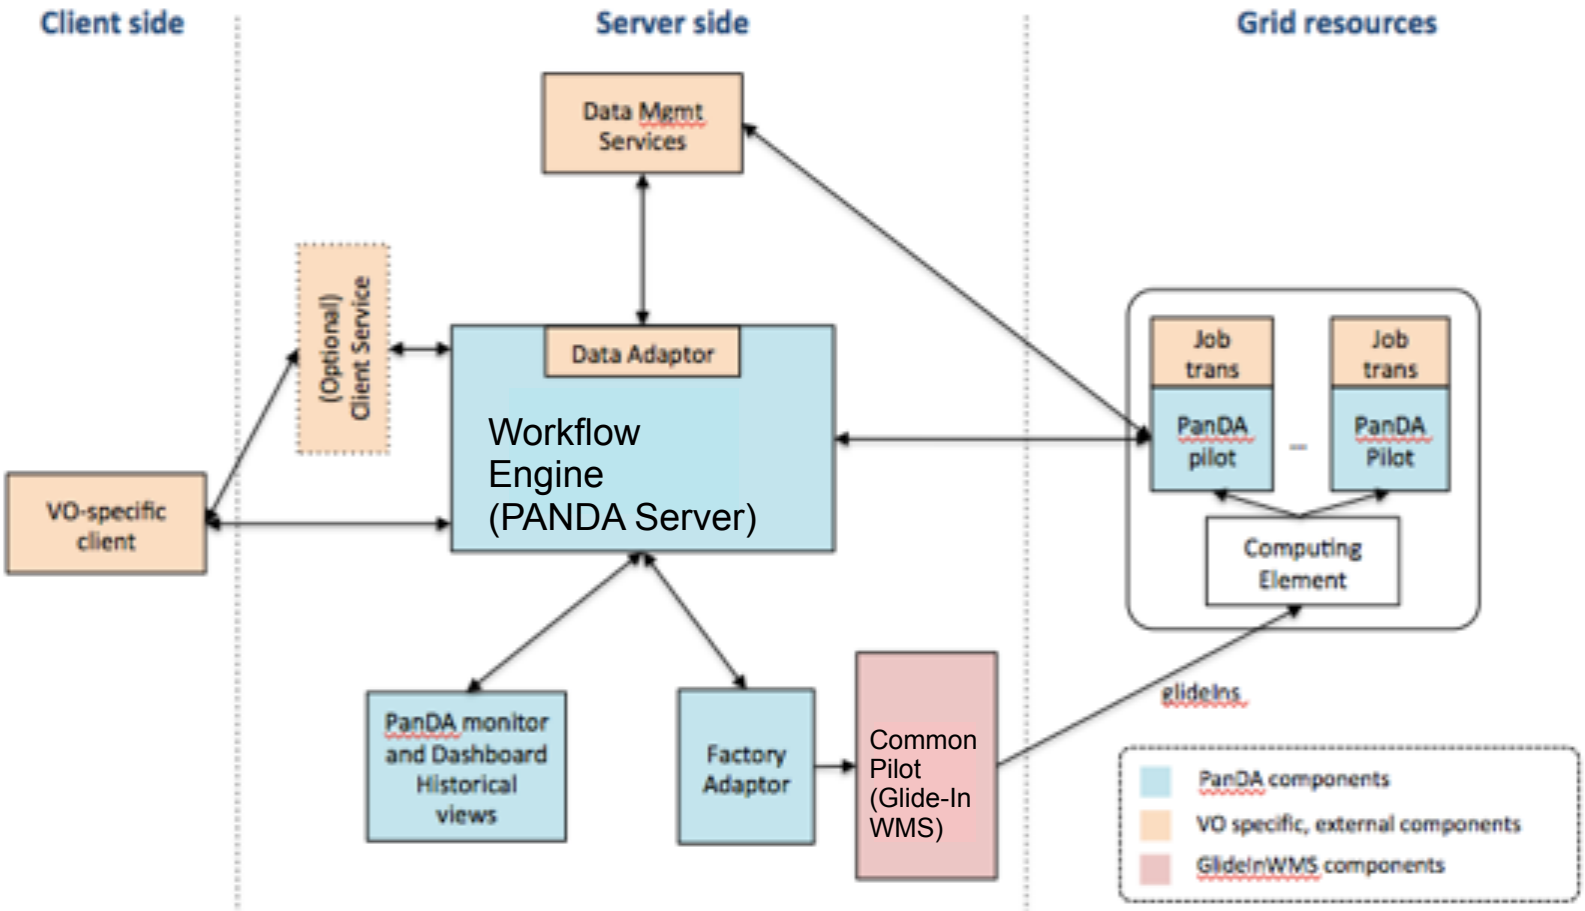


Ian Fisk
FNAL/CD

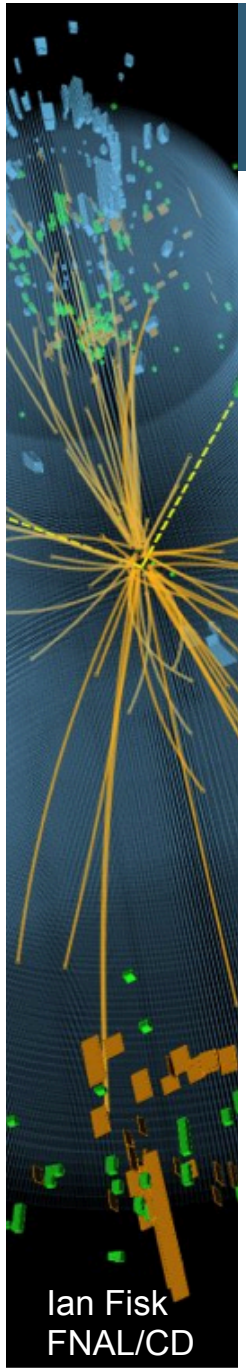


# Plan

- Completed the Feasibility Study in May
  - The component functionality and interactions were examined and no show stoppers were identified to exploring common prototypes
  - Pursuing a 6 Step approach for a Proof-of-concept Prototype
    - STEP 0: Run Basic CMSSW job
    - STEP 1: Include pilot factory
    - STEP 2: CMS client tool
    - STEP 3: Output file handling
    - STEP 4: CMS output management
    - STEP 5: Log and output access
  - Goal is to have a functional prototype by the fall to decide to be able to make informed decisions about moving on a common product

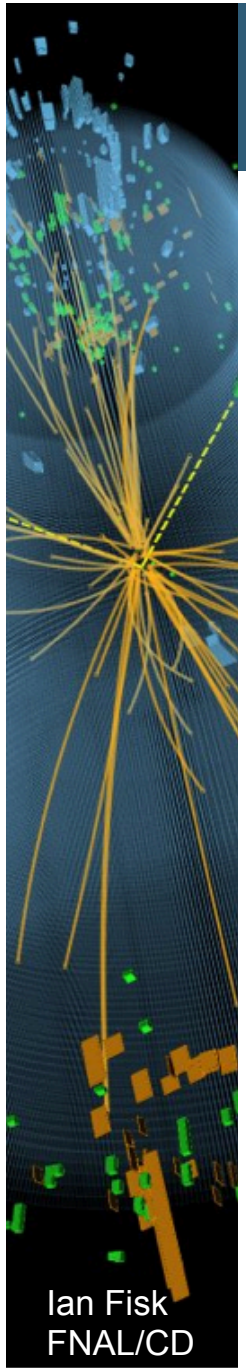


# Progress

- We have reached the level that CMS has been able to submit basic jobs to the PANDA server
  - A lot of the code for job specification and splitting from the current system can be re-used
  - Involves some reorganization of code to make it more experiment generic
- We want to connect the production PANDA server from ATLAS with the production Glide-in WMS system from CMS to demonstrate the scale possible
  - Involves some development to reasonably handle resource allocation across two experiments

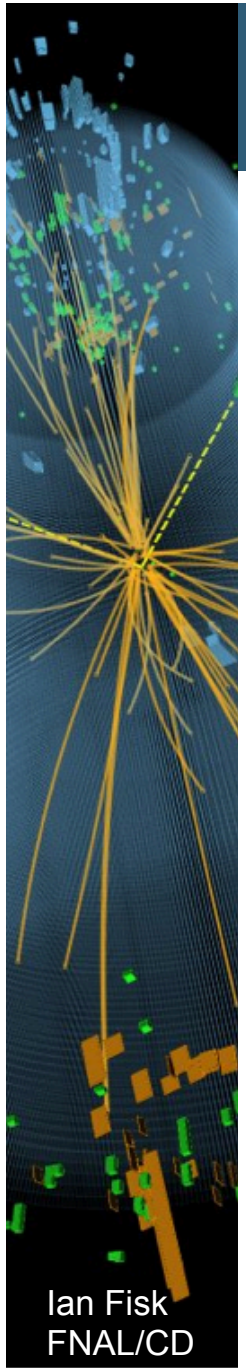


# Outlook

- IT-ES/VOS has a good record of identifying and development common solutions between the LHC experiments
  - Setup and expertise of the group have helped
- Experiments are engaged and interested in the process and the development
- Several services focused primarily on monitoring have been developed and are in production use
- More ambitious services that would be closer to the experiment core workflows are under investigation