
GridPP
UK Computing for Particle Physics
iris
Activities using
GridPP @ IRIS
Week 09
Week 10
Week 11
CPUs
Running Processes
last month
IRIS – Cambridge
14-Jan-2026
Pete Clarke

Intro statements

- All (almost all ?) large-scale Particle Physics(PP) experiments use the global PP Grid for computing
- Some other large- scale experiments (LZ and Clas12) are very similar in the way they compute and so also use the global PP Grid infrastructure
- IRIS supports this by supporting GridPP

What is GridPP ?

- GridPP provides computing for all the LHC experiments.
- GridPP also provides computing for essentially all other HEP experiments funded by STFC (that are large enough to need distributed computing).
  - This is at marginal costs of staff-effort and hardware.
  - The small ones (~ 1-2% of GridPP capacity) are covered under GridPP
  - The larger ones obtain their capacity allocations via IRIS.
- Current VOs using GridPP @ IRIS include
  - Dune
  - LZ
  - Microboone, SBND
  - Clas12
  - Darkside
  - LIGO
  - LSST
  - Euclid

GridPP
UK Computing for Particle Physics

What is GridPP?

- GridPP is a leading partner in the Worldwide LHC computing Grid (WLCG):
  - WLCG is 40+ countries, 164 sites, bound together by 64 MoUs.
  - Tier-0 @ CERN; ~12 Tier-1; 150+ Tier-2 sites.
  - Global networking up to ~500 Gbit/s.
  - **1+ Million cores and several Exabytes of data.**
  - **GridPP is 13.5% of the Tier-1/Tier-2 infrastructure of WLCG**
  - Worldwide context brings constraints and responsibilities: Long established international, standards, agreements, authorisation and security requirements.

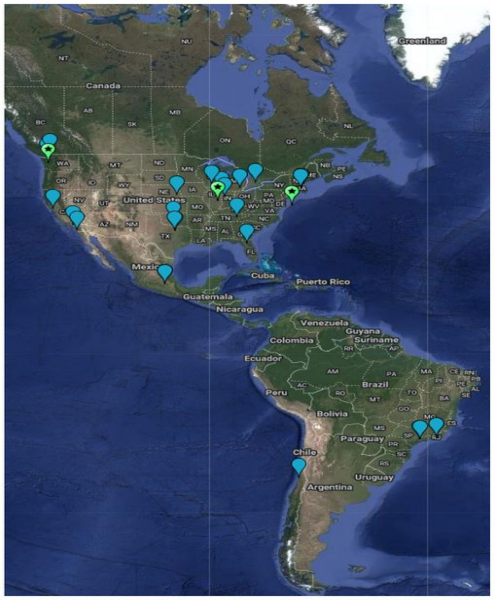

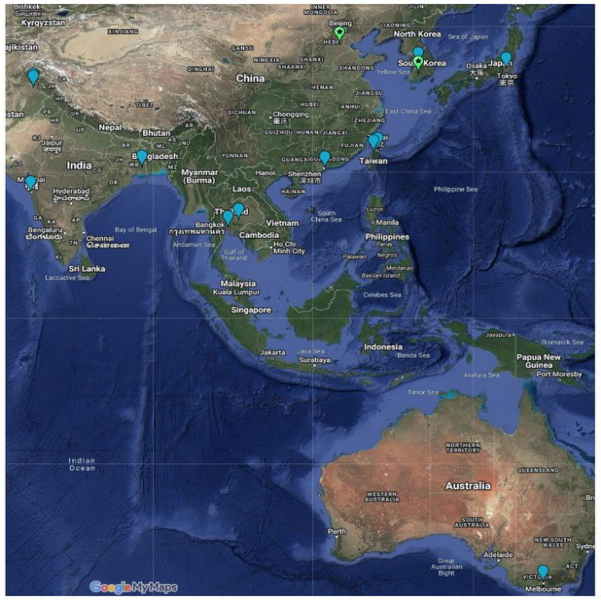

What is GridPP?

- GridPP started 25 years ago:
  - Tier-1 at RAL (in WLCG parlance)
  - 5 large core Tier-2 sites
  - 11 smaller Tier-2 sites
  - ~120,000 compute cores
  - ~110 PB of disk
  - ~150 PB tape
  - 100 Gbit/s networks

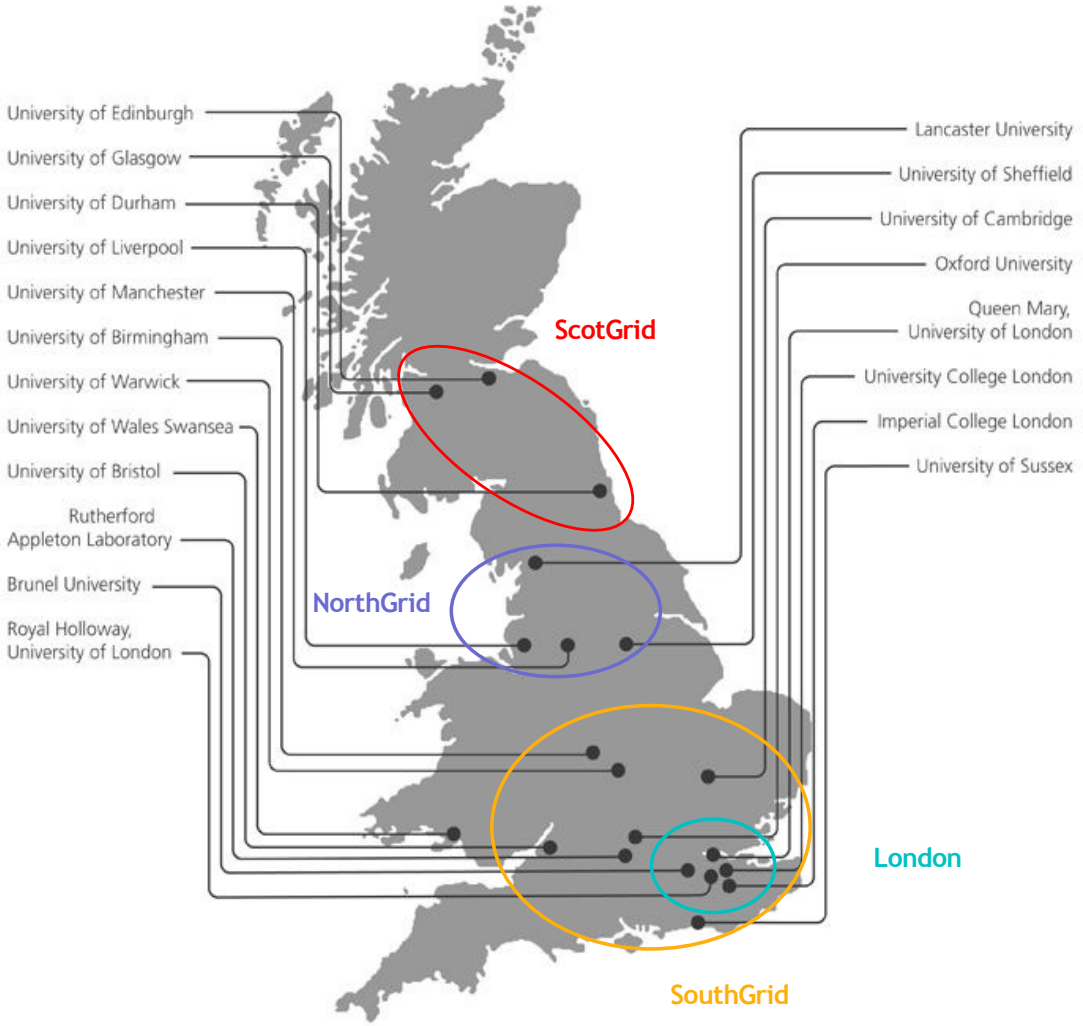

Much more than hardware!

- Also a comprehensive software fabric that enables seamless worldwide distributed computing (green layers).
- This allows the LHC experiment teams (orange layers) to:
  - Use the WLCG as a seamless cluster of ~ 1 million processors.
  - Replicate, locate and use data anywhere around the world (mostly Rucio)
  - Automatically send millions of compute jobs around the world to the data and retrieve the results.
- If all hardware (blue) were provided free by someone else, all the software fabric (green), and the people to run it, would still be needed (none of this is provided by any cloud service, for example).

Layer	Responsible	Experiment 1	Exp. 2	Exp. 3
6	Experiment physicist end users	Selecting data, running analysis code.	…	…
5	Experiment physics programmers and software engineers	Analysis frameworks, reconstruction code, calibration code…	…	…
4	Experiment computing teams	'Production' computing operations and software	…	…
3	WLCG/GridPP	Middleware interface to experiments, and experiment 'customer' support		
2	WLCG/GridPP	Software infrastructure running on physical hardware infrastructure and WLCG Federal responsibilities		
1	WLCG/GridPP	Physical Hardware		

Illustration of complexity and pervasiveness of operating fabric

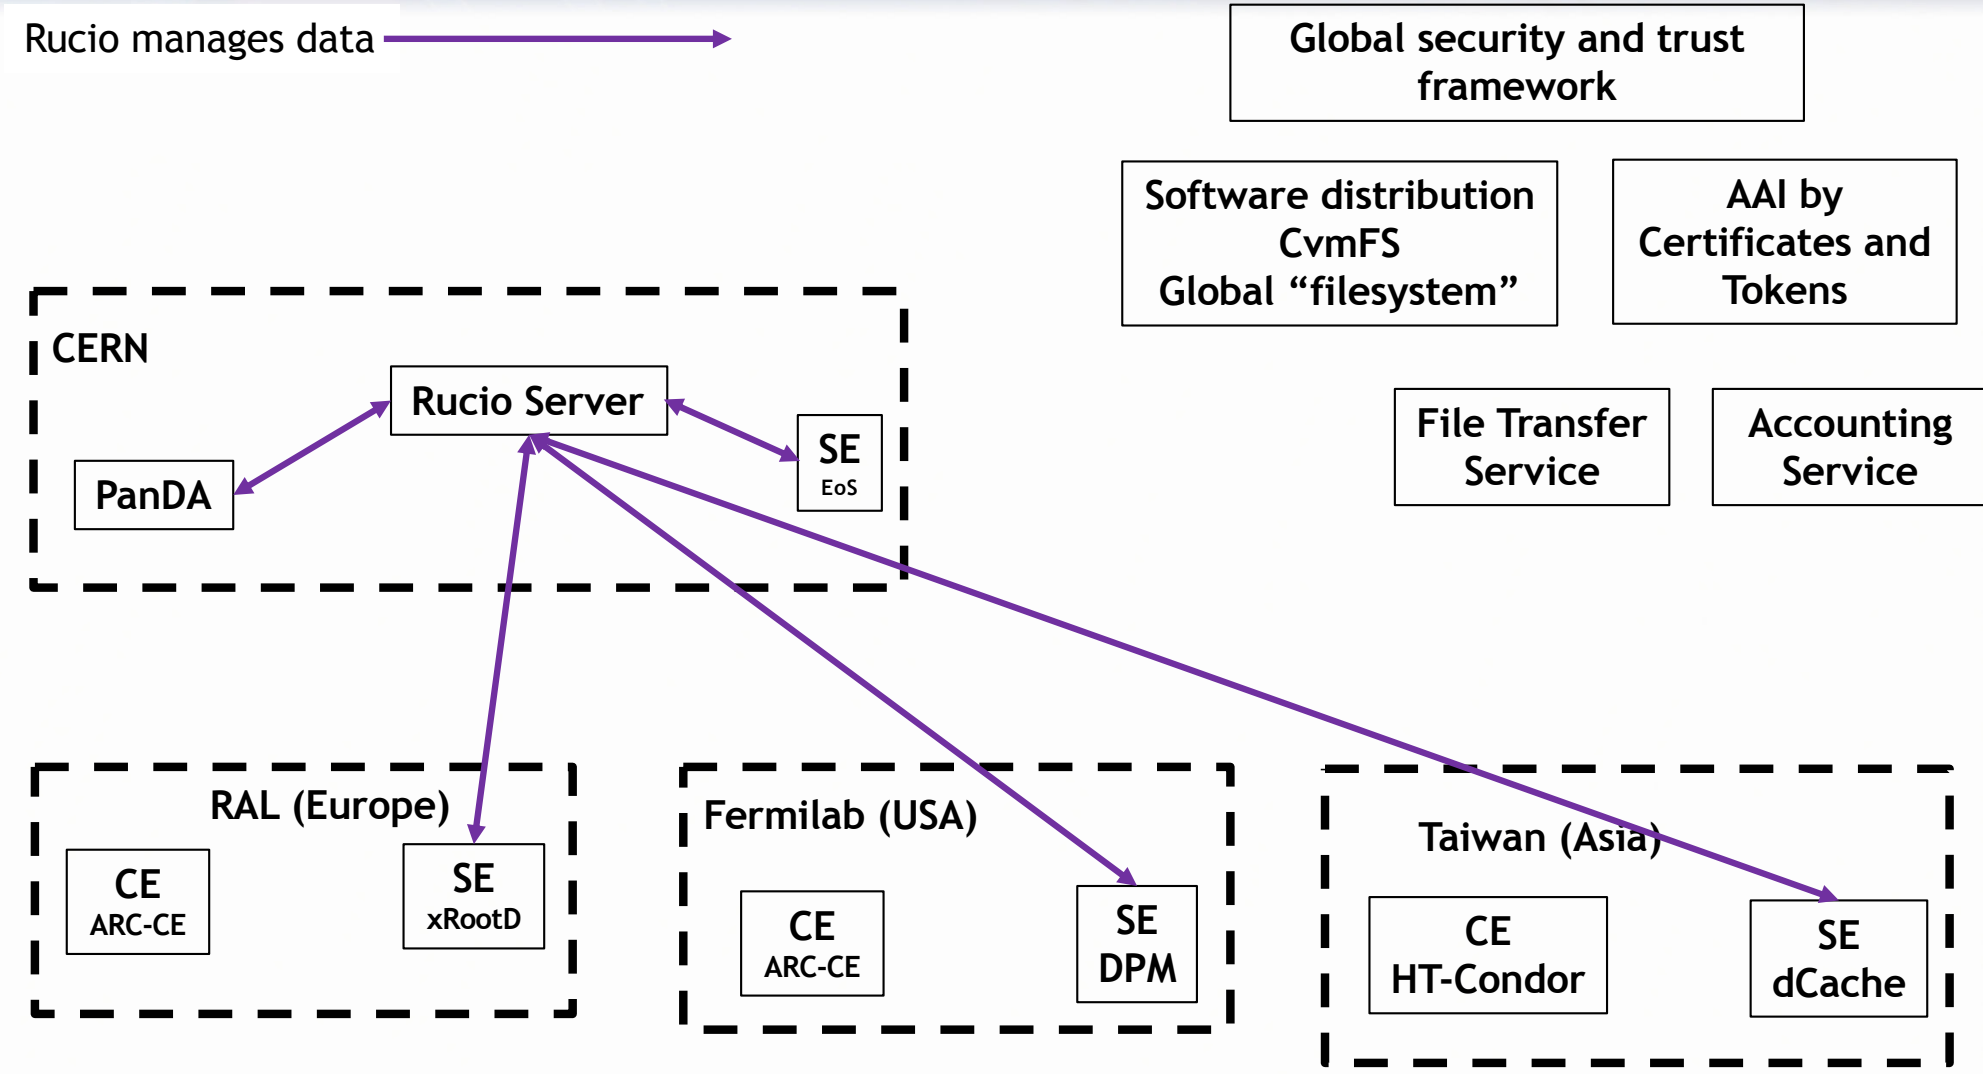

Illustration of complexity and pervasiveness of operating fabric

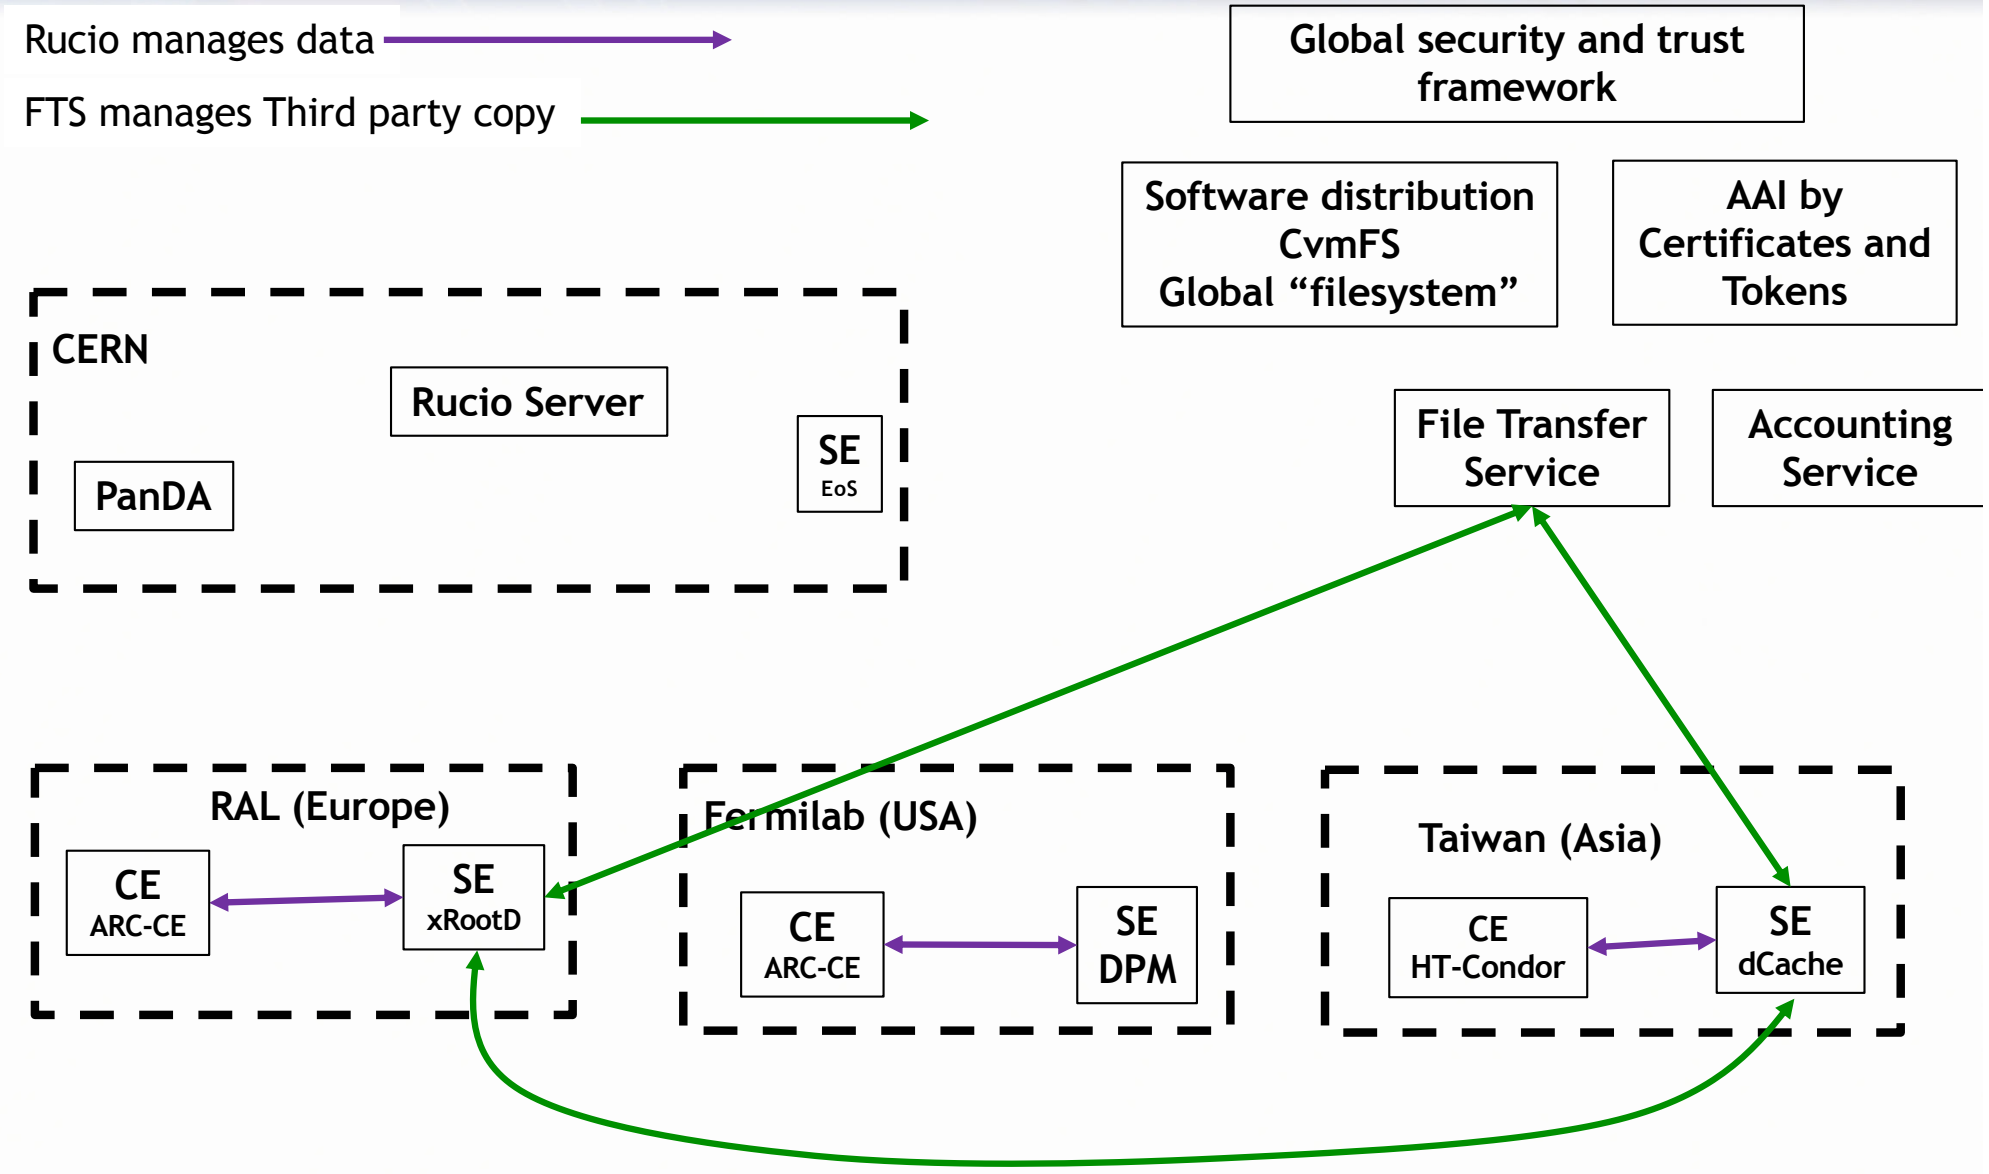

Illustration of complexity and pervasiveness of operating fabric

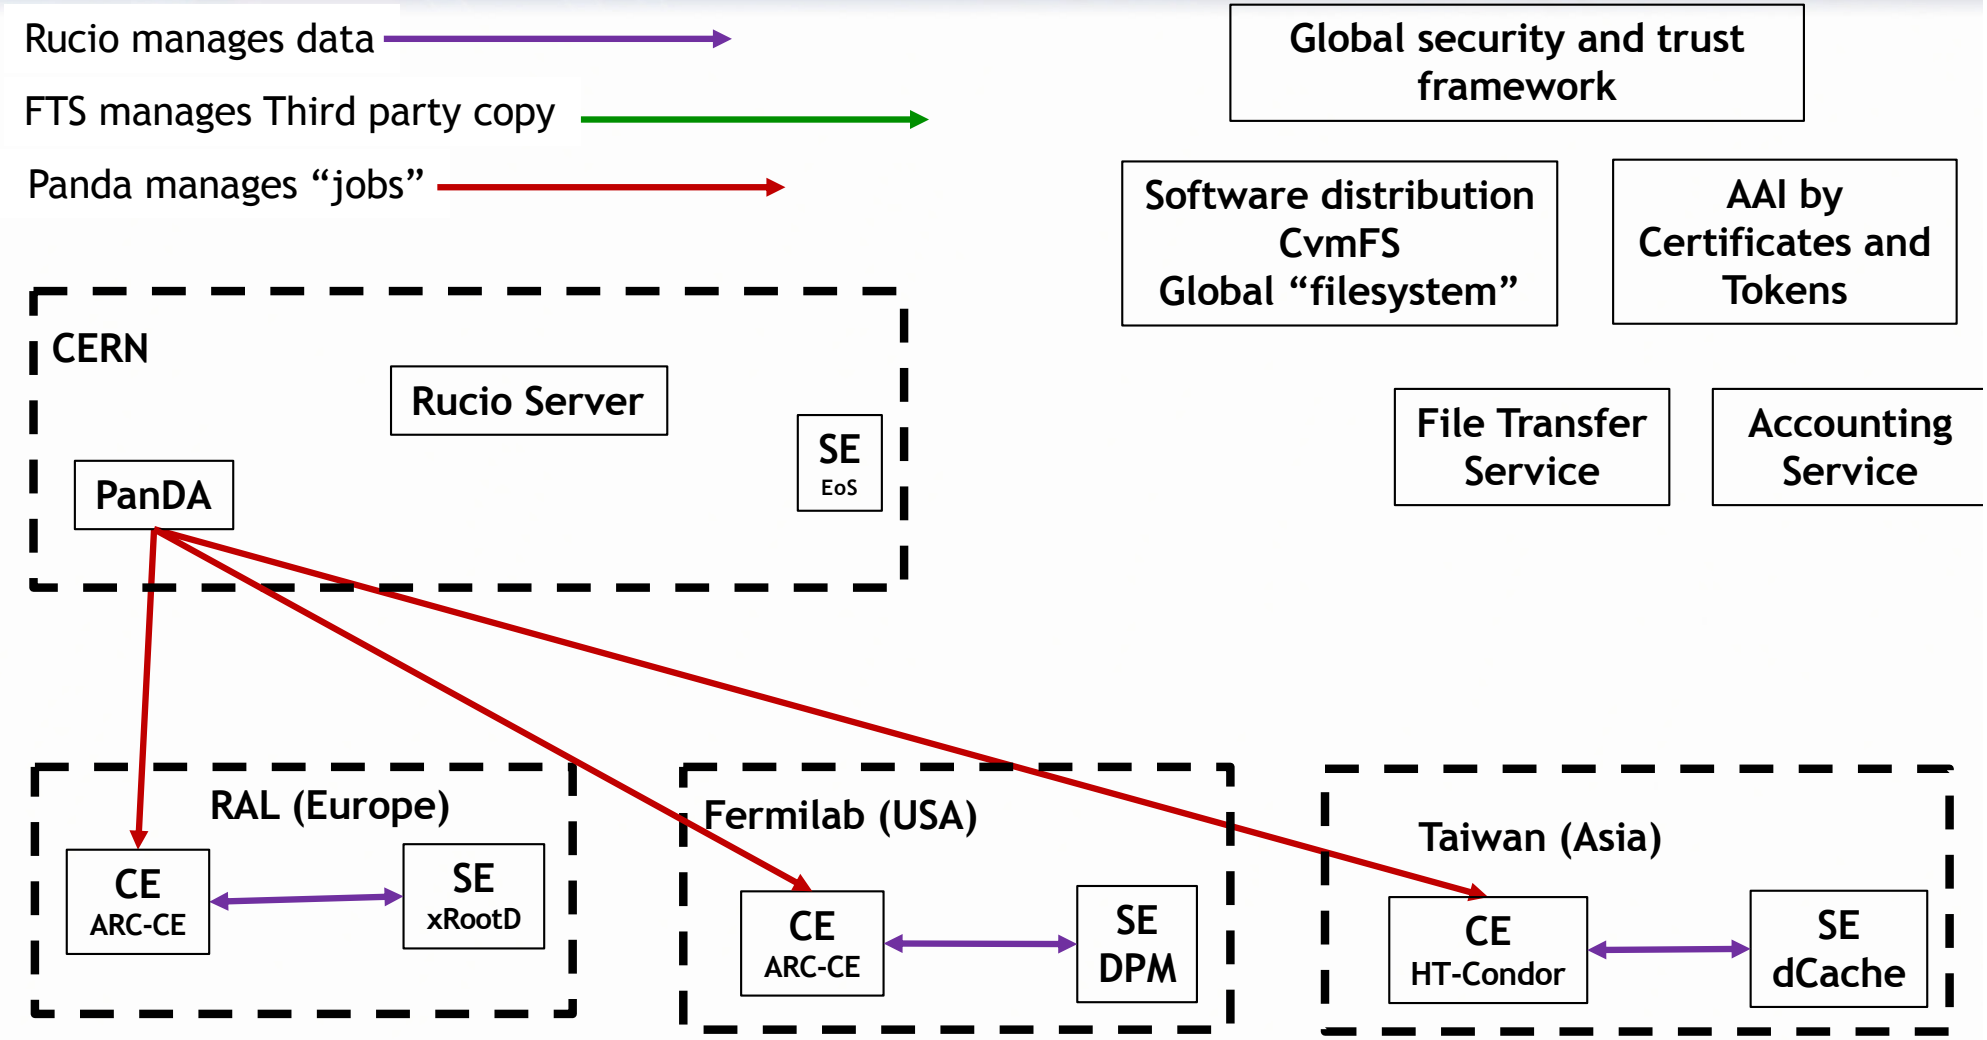

Illustration of complexity and pervasiveness of operating fabric

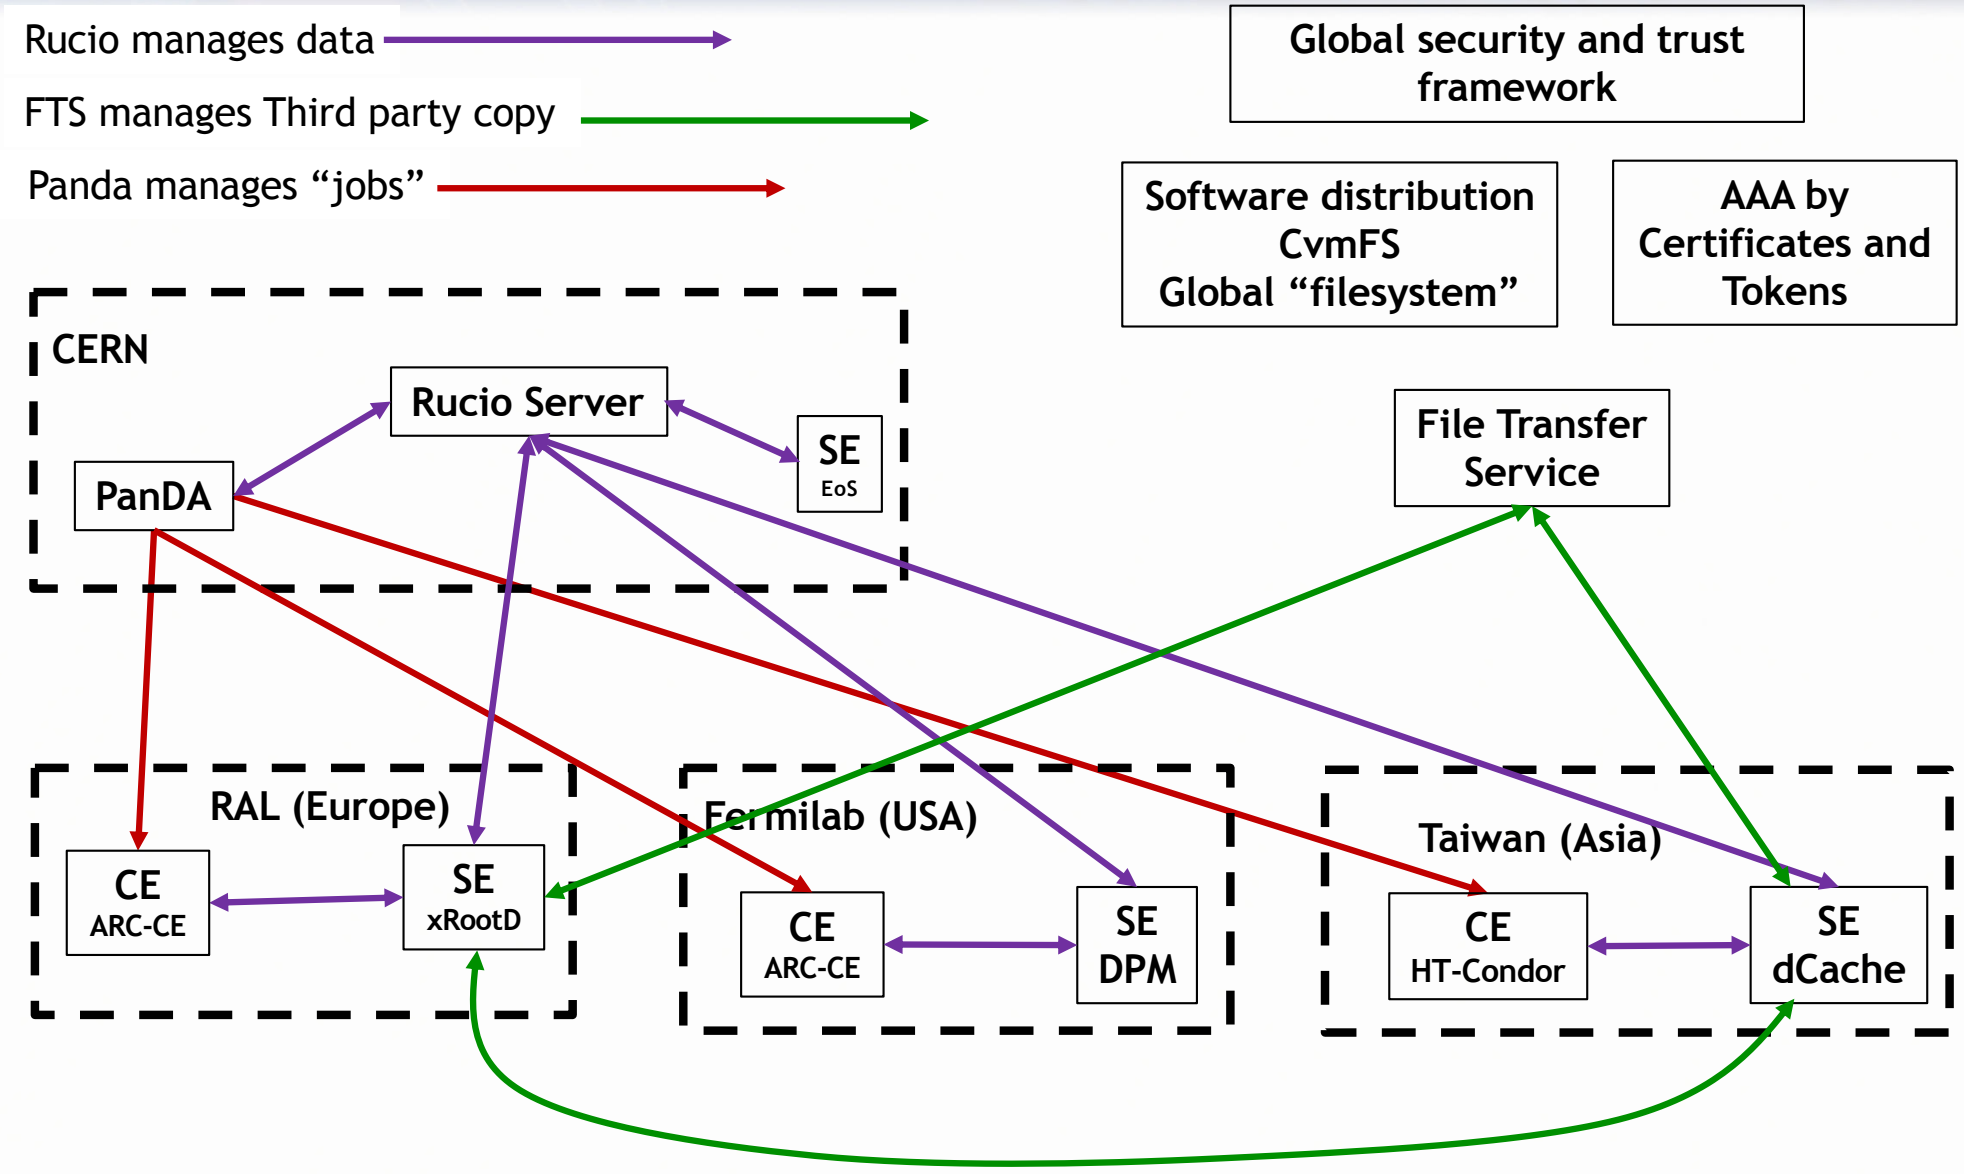

Activity talks from:

- The point being that all Activities using IRIS@GridPP use this underlying infrastructure with their particular Computing Model and “orange software layers”

- Now slides from activities themselves…

Questions:

- What has IRIS enabled you to achieve ?

Zeroth order answer is: Allow X to work

- Allow UK to provide its fair share of deterministic compute for X – and this alone is a first-class achievement.
- It is X itself that pushes back frontiers of science and publishes papers

- What is top of your wish list for computing (that IRIS could help with) ?

Zeroth order answer is: Provide/lobby for software engineering effort

- What could IRIS do to help new “small groups” use the Grid infrastructure ?