Software & Data Infrastructure for Earth System Modelling

Bryan Lawrence

STFC Centre for Environmental Data Archival

(NCAS British Atmospheric Data Centre

& NCEO NERC Earth Observation Data Centre

& IPCC Data Distribution Centre)

Eric Guilyardi

IPSL &

NCAS Climate

Sophie Valcke

Cerfacs

British Atmospheric Data Centre
NATIONAL CENTRE FOR ATMOSPHERIC SCIENCE
NATURAL ENVIRONMENT RESEARCH COUNCIL

Institut Pierre Simon Laplace

National Centre for Earth Observation
NATURAL ENVIRONMENT RESEARCH COUNCIL

Outline

- A philosophy for a strategy for integrated software support for Earth System (Modelling) Science.
- Existing practical endeavours to support integrated descriptions of simulations, and integrated data management
  - CMIP5 = Simulations + ESG (access) + Metafor (&CF) (provenance) + Local Solutions (analysis)
- Possible futures for software infrastructures
  - From frameworks to portals
- Summary: Networking and self-governance (or directed governance) just as important as the software.

Today → Tomorrow

- Models developed *independently* and integrated (sometimes) in *parochial* frameworks with various level of *support* for their usage communities.
- Data held in local archives (sometimes nationally), with IS-ENES working on distributed database concepts, but poor *distributed access.*
- Good and improving support for CF-netcdf in the modelling community + immature model descriptions from Metafor.
- Data manipulation and access tools developed *independently* and with various levels of *support.*

→

- Models developed by communities working with common coding conventions and shared support.
- Data held in distributed archives, with key data sets aggregated and replicated as necessary, with well understood routes for moving data as necessary.
- CF continues to be supported, but now prevalent in the EO and observational communities. The information and vocabularies built by Metafor are maintained by an international community.
- Data manipulation and access tools are developed by communities working with common conventions and shared support.

Problems solved along the way

- Securing common goals (without which the rest wont matter).
- Resolving the contention between supporting both innovation and efficiency of distributed development.
- Establishing common software conventions.
- Putting in place governance to allow the conventions to evolve.
- Moving communities to understand that runtime is not the only thing to optimise: development time needs optimisation too.
- Procuring dedicated network paths ("light paths") along major data routes.
- Changing working practices to ensure that simulation data is
  - appropriately (in physical space) stored, and
  - documented well enough to be reused (or discarded) appropriately.

Expected benefits

High performance Community ESM with infrastructure

- developed by dedicated experts, available to institutes/teams at low effective costs
  - Helps scientists to focus on science.
  - Helps scientific diversity
    - Survival of smaller groups
    - Range of types of ESM easier to develop and support (e.g. from regional predictions to paleo climate)
    - Community software implies more scientific exchanges
  - Encourages computer manufacturers to contribute to
    - efficiency (porting, optimisation) on variety of platforms
    - migration effort for next generation HPC
  - Reduced overall costs
    - easier procurements, migration, and benchmarking

CMIP5, an exemplar: The software runs through it!

The Fifth Coupled Model Intercomparison Project

- Sponsored by the WMO WGCM
- Quality Controlled Data to (eventually) appear in the IPCC Data Distribution Centre
- Data pipeline starts as a WGCM issue, morphs into something loosely ~~governed~~ organised by GO-ESSP, and then become (in part) an IPCC-DDC issue …

• Every European modelling project will/should have similar software requirements in the pipeline – from producers to consumers...

CMIP5 & ESG: A more European view

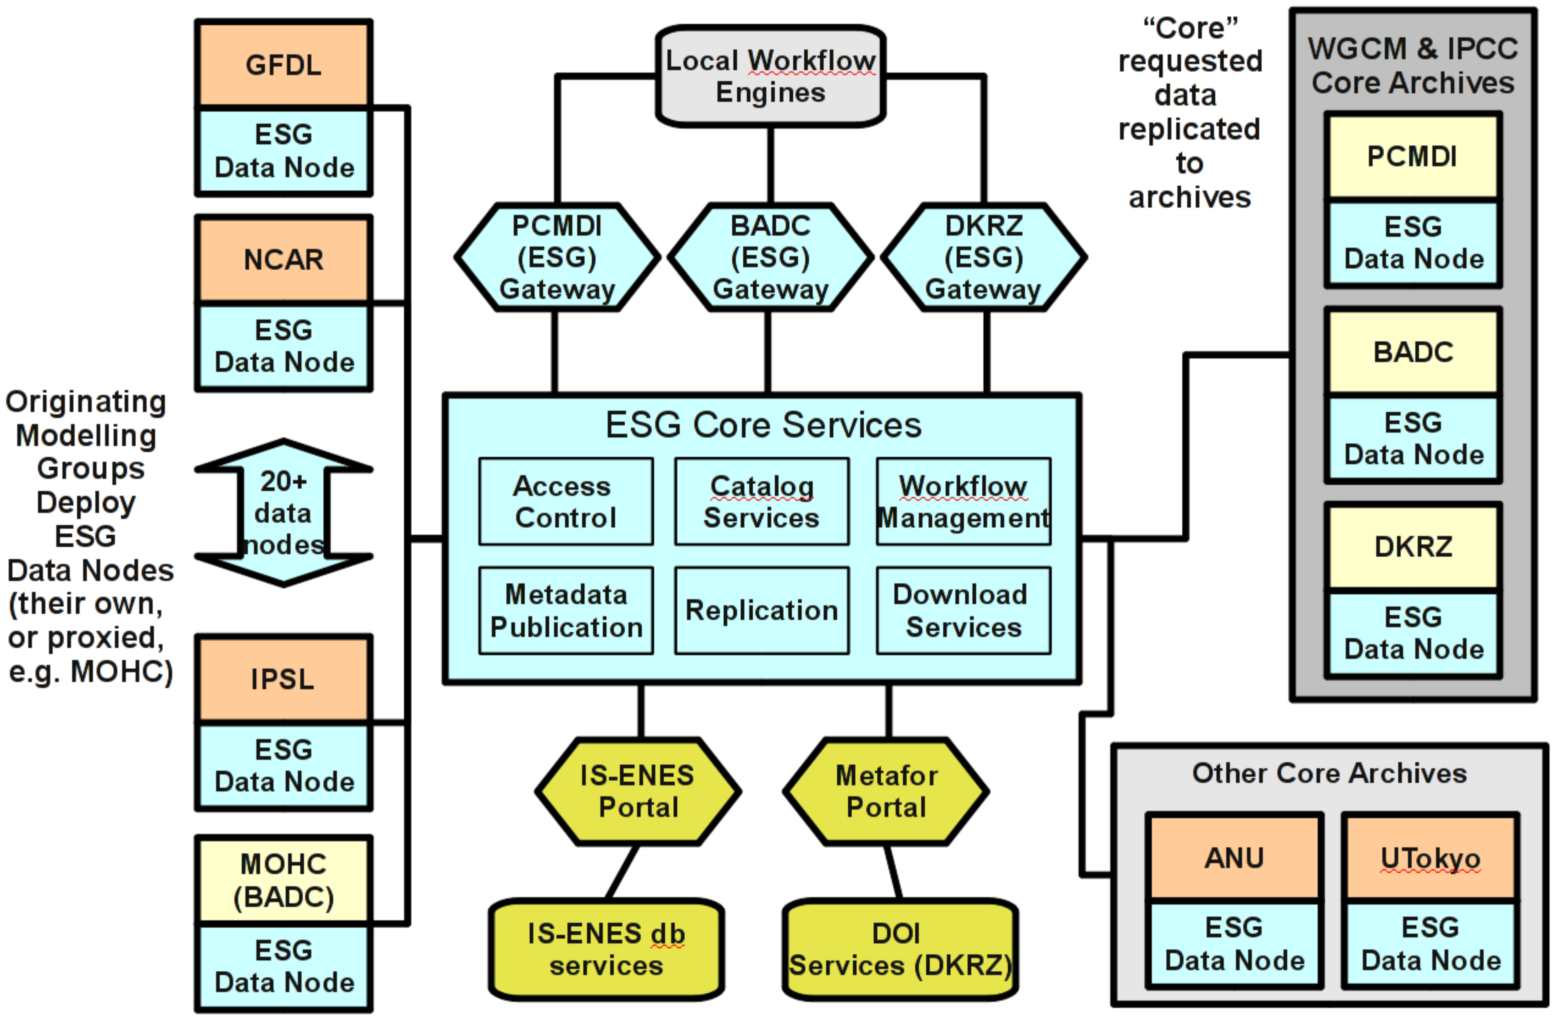

Managing Data: CMIP5 information view

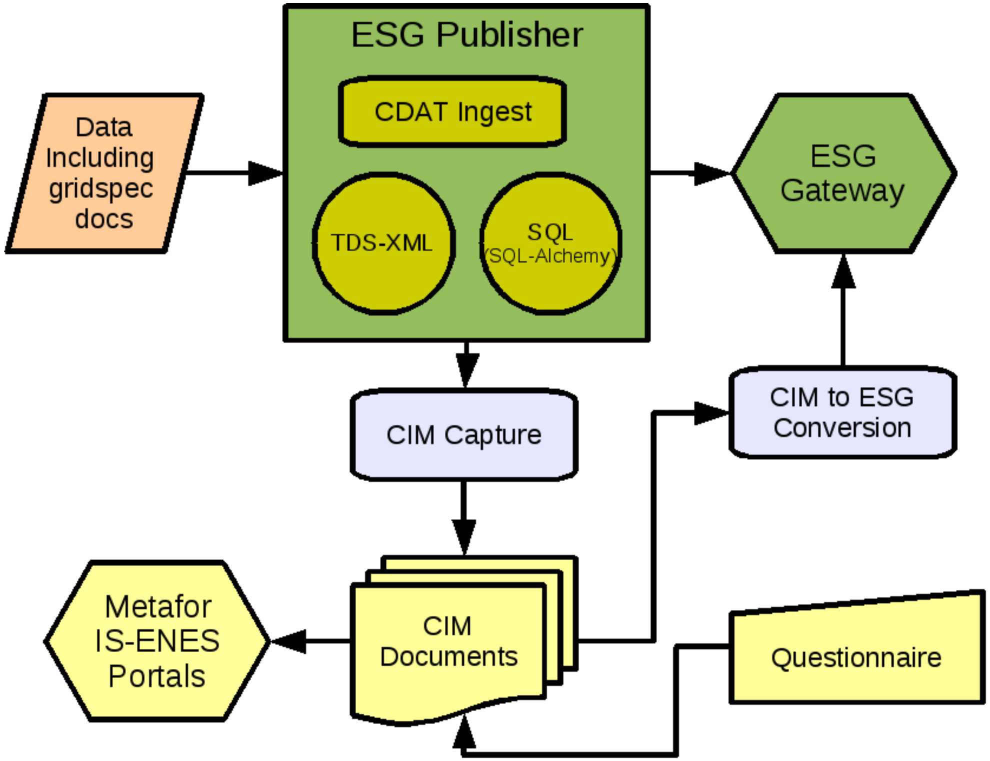

Provenance comes via two paths:

- CF file attributes and
- Metadata Questionnaire

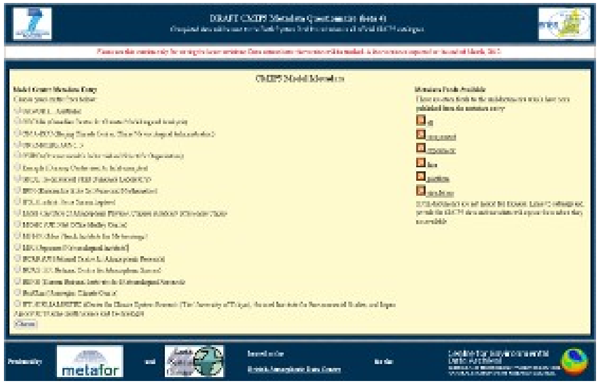

Stored description documents served up by services used in multiple portals (at least: ESG Gateways, Metafor Portal and IS-ENES portal).

Managing Data: CMIP5 in numbers

Simulations:

~90,000 years

~60 experiments

~20 modelling centres using

~30 major(*) model configurations

~2 million output datasets

~10's of petabytes of output

~2 petabytes of CMIP5 requested output

~1 petabyte of CMIP5 "replicated" output

- Which will be replicated at BADC & DKRZ, to arrive in 2010/2011!

~10 TB of land-biochemistry (from the long term experiments alone).

Of the replicants:

~ 220 TB decadal

~ 540 TB long term

~ 220 TB atmos-only

~100 TB of 3hourly atmos data!

~215 TB of ocean 3d monthly data!

~250 TB for the cloud feedbacks!

Expected Usage (@ BADC):

~ hundreds of users downloading at a sustained daily average rate in excess of 1 Gbit/s (up to 35 TB/day from BADC ...)

The pieces of CMIP5 support > £2M!

HARDWARE
- Data storage
- Approx. 1,000TB
- Faster network
- Servers to deliver and process data

SOFTWARE COLLABORATION
- International effort
- Replication system
- QC & Versioning systems
- Harmonisation
- Interfaces to data

METADATA DEVELOPMENTS
- International effort
- Describing models, experiments and datasets
- Standard format and description for all

USAGE TOOLKITS
- Sub-setting
- Batch processing
- Re-gridding
- Format conversion
- Visualisation
- Analysis Platform

MOHC DATA SUPPORT
- Data handling of MOHC models
- Checking and QC
- Connection to tools

NERC DATA SUPPORT
- Data handling of HIGEM and Paleo models
- Format conversion
- Checking and QC
- Connection to tools

INTERNATIONAL DATA SUPPORT
- Data handling of models
- Checking and QC
- Connection to tools

UK Community
- Engagement with Impacts Community
- Public Sector, general public and Private Sector access
- Development of Derived Products

(From a BADC perspective). Not just about software. Indeed, not even mainly about software!

ESM data in the wider community

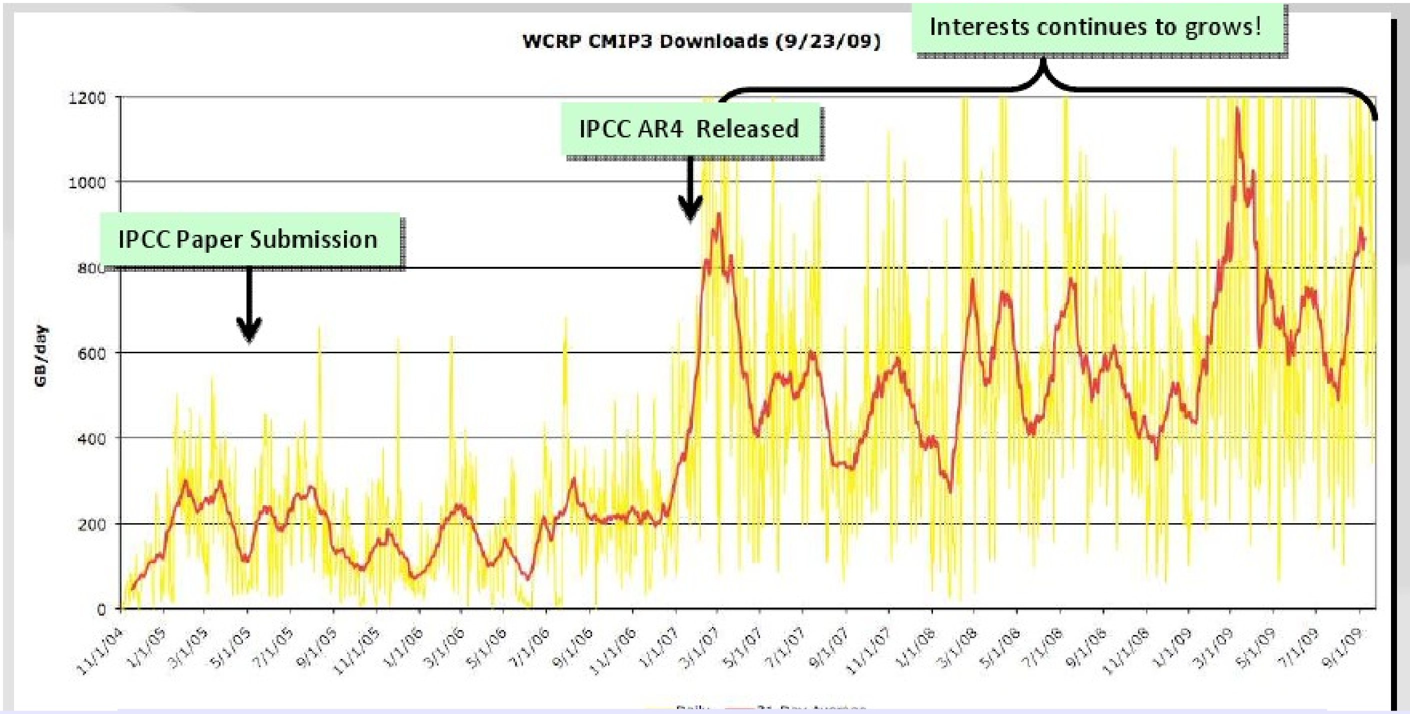

Courtesy of Bob Drach and Dean Williams PCMDI

Huge interest and need for ESM data in the wider group beyond the traditional ESM community.
This community is the reason we exist!
We need to do better at supporting them = Support them with familiar software!

…and so to building models:

Software structure of an Earth System Model

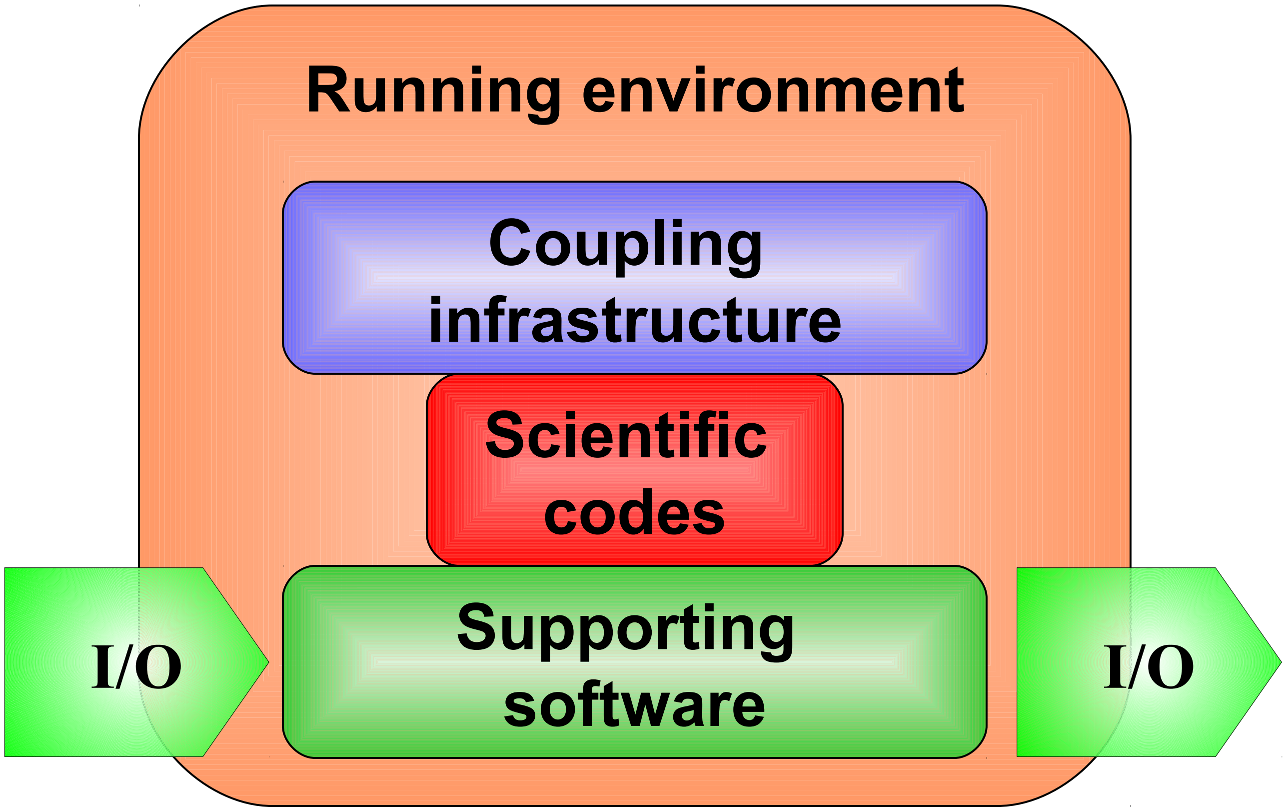

ESM Software Development

Success factors (left) and distinguishing traits (right)

Success factors
Tight integration between science and coding.
Single site development (of major components) *(is this really true, or about who was evaluated?)*
Software developers are domain experts.
Shared ownership and commitment to quality.
Openness
Benchmarking (model intercomparison)
(Nearly) unconstrained release process – new models only released when "good enough"

Easterbrook & Johns (2009)
Comp. Sci. Eng.
DOI:10.1109/MCSE.2009.193

Distinguishing traits
A stable architecture (in the sense of being based on the physical components)
Modules and integrated system owned by different communities
The programming language for modelling is Fortran & the people are smart
Software has huge societal importance but individual bugs unlikely to lead to radical changes in results.
Existence and use of coupling frameworks (but everyone wants better than they have)
Testing focusses on integration (not unit tests)
Few resources for software infrastructure

Easterbrook et al (2010)
Blog: http://www.easterbrook.ca/steve/?p=1558

(not yet clear whether these results will stand comparison with lots of other ESM groups)

Coupling Technologies for ESM

Different technical solutions are used in the ESM community to couple geophysical model codes

Two main approaches (besides hard-coding):

- Use an external entity (for transformation of the coupling fields) and link its communication library to *existing* applications sometimes referred to as "coupler" approach in the community
  - e.g. OASIS
- Use coupling library/functions to *build* an integrated coupled application based on elemental *science units*, sometimes referred to as "framework" approach in the community.
  - e.g. ESMF

The different implementations of coupled models in the community lie in the continuum between those two approaches.

Coupling Technologies for ESM: The "coupler" approach

- Keep original codes almost unchanged and interface them with a communication library
- Use an external "coupler", configured by the user, to transform the coupling fields

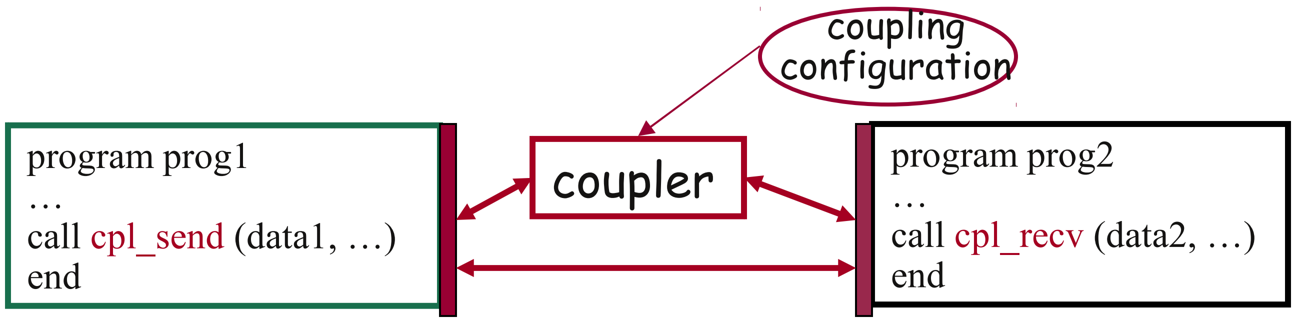

☺ change existing codes as little as possible
☺ flexible
☺ portable
☺ use of generic transformations/regridding
☺ concurrent coupling

☹ efficient
☹ sequential coupling

→ probably best solution to couple independently developed codes

Coupling technologies for ESM: the Framework Approach

- Split original code into elemental units
- Write or use coupling units
- Use the library to build a hierarchical merged code
- Adapt code data structure and calling interface

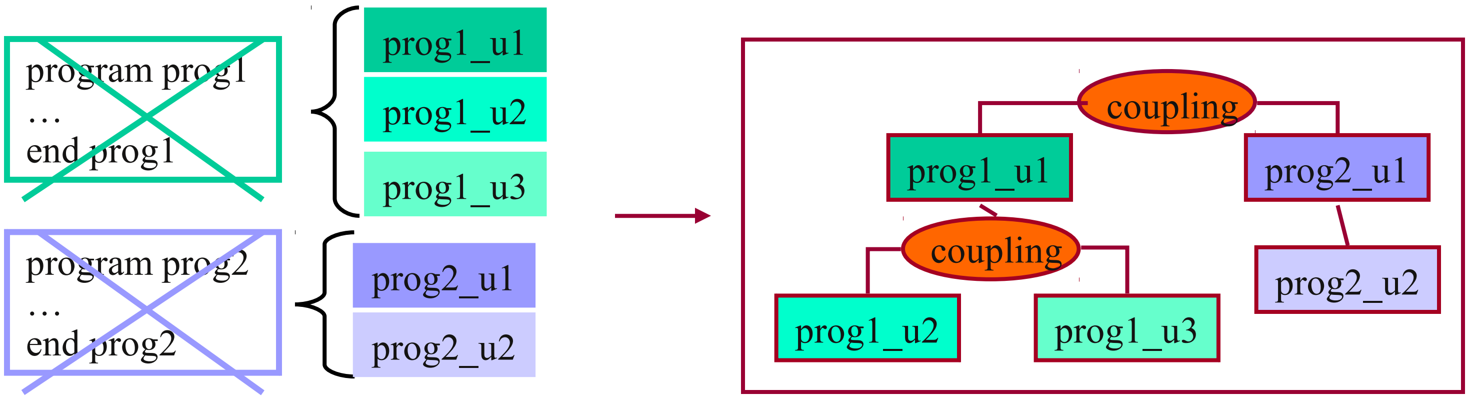

☺ efficient
☺ flexible
☺ portable
☺ use of generic utilities (parallelisation, regridding, time management, etc.)
☺ sequential and concurrent components

☹ existing codes have to be modified
☹ not easy to start with

→ probably best solution in a controlled development environment

Coupling Technologies for ESM: The Future.

The "best" coupling technology does not uniquely exist; it depends on:

- the level of change/adaptation to existing codes one can support
- the efficiency one wants to achieve
- the computing environment
- the required utilities
- the level of agreement to conventions that can be achieved

In Europe, given the diversity of the developing institutions, most of the groups naturally adopted the "coupler" approach with OASIS:

- is this the right way to go on?
- are we ready to agree on, adhere, and resource more conventions and constraints (coding rules, etc.) to be able to build more integrated and more efficient coupled applications?
- will we be forced to, given the future computing platforms?

… but it's not all about the model code!

- Common approaches to Inputs and Outputs require:
  - Common file formats (netCDF), and
  - Common file conventions (CF-netCDF).
    - ☺ CF prevalent in the community, but not yet the default I/O format for models, leads to unnecessary file manipulation …
- Common approaches to analysis require
  - Correlative data (“observations” and “earth observation”) to conform to the same file formats and conventions.
    - ☹ Much needs to be done to help CF generally for observations and EO and the promulgation of netCDF into EO.

Using more computing: and the common factor is?

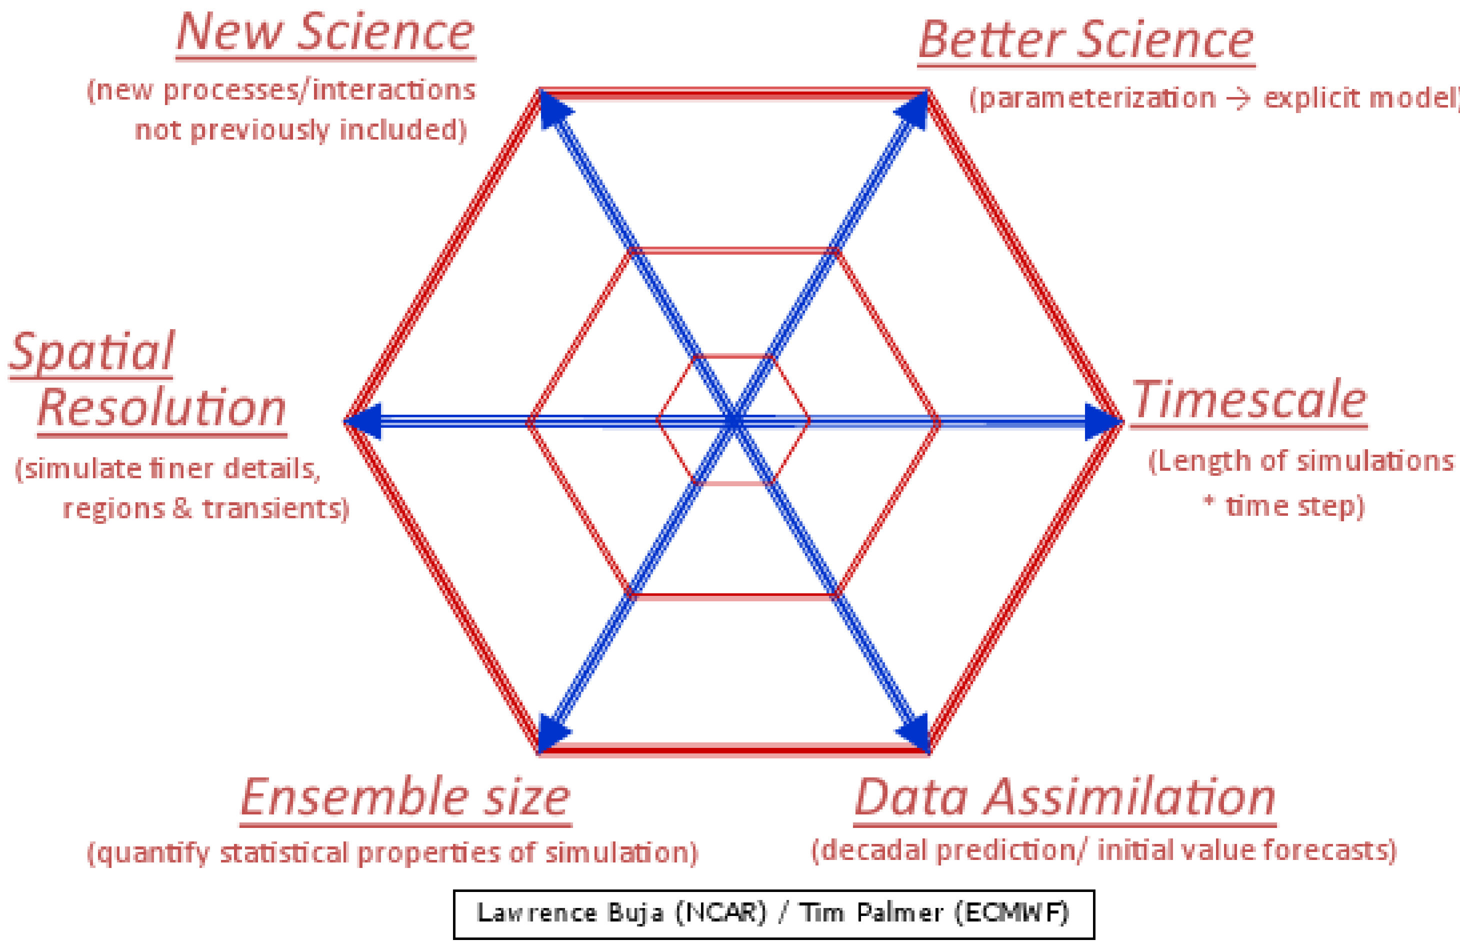

Using more computing: and the common factor is?

Data handling & model Development & more correlative data

Data handling & model development

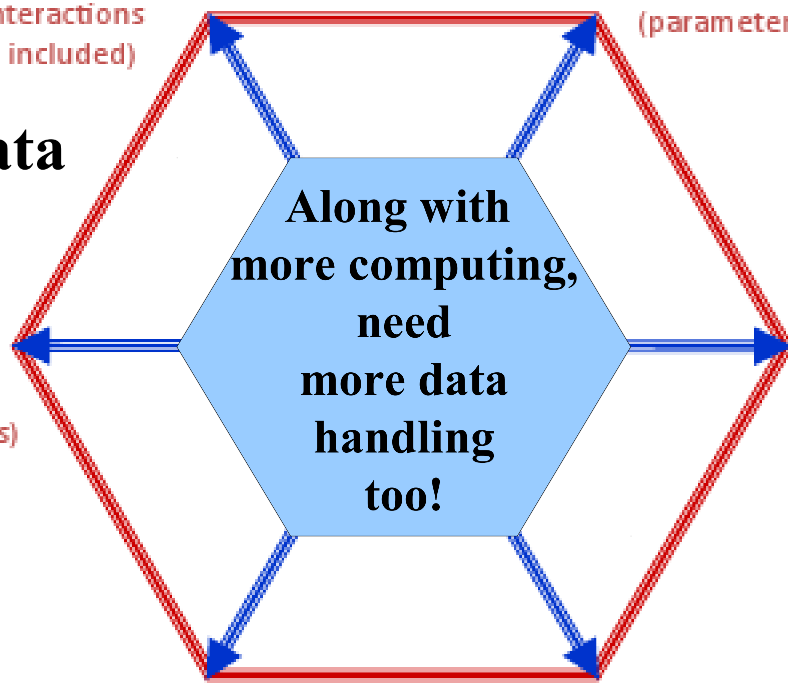

Data handling & better science & timestep & more correlative data

Data handling

Data handling & analysis tools

Data handling & mathematics & input data

Post-processing matters too!

Most post-processing and visualisation tasks using lengthy simulations end up spending most of their time reading and writing data (I/O), and relatively little doing calculations.

- That might not be true if folk used high temporal or spatial resolution for analysis, but that's rare …
- … many argue using full resolution data for analysis is unnecessary, often from habit rather than logic (although sometimes it's true!)
- … many of those same folk plead for higher resolution modelling (but are fixated only on the upscaled linearly averaged effects).

What if it was much easier to

- Move data? Compare Data? Aggregate non-linear high-resolution calculations, rather than average first?
- Do non-linear calculations at high resolution and calculate higher order statistics directly?

Do we do too much data analysis on expensive super-computer hardware because that's where the fast disk is, and it's too time-consuming to move the data?

Using more computing: conclusions

- Not all options are served by one big machine, which implies
- Multiple machines, and multiple locations where codes are integrated, which implies
- Post-processing (differencing etc) requires bringing data together across networks, but
- If we centralised, we'd have to backup, and get no benefit of having those mutiple data copies, and
- Not all data is HPC data, so location of correlative (EO etc) data matters too, which implies
- We can't centralise (on a European scale) data analysis,

So we need distributed (& more sophisticated) solutions for data analysis as we exploit more computing capacity

Moving data analysis along ...

We need to

- Invest in the software infrastructure for data analysis
  - There are islands of good practice (CDO, NCO etc along with NCL, CDAT, GraDS and others)
  - All based around a single institution (?)
  - What is the right European approach – to support both efficiency and competition/diversity?
- Address server-side processing, GRID based data analysis & private clouds alongside managed data archives.
  - In the limit of n=many HPC sites, all data will move at least once, but we don't want all data moving n times.
  - Solution: m managed data archives (m<n), with fast disk, and analysis clusters deploying virtual clouds &/or server side computing (GRIDs)
  - On a national scale m=1 or 2 is probably appropriate!

Modelling Post-Processing Infrastructure: UK-centric View

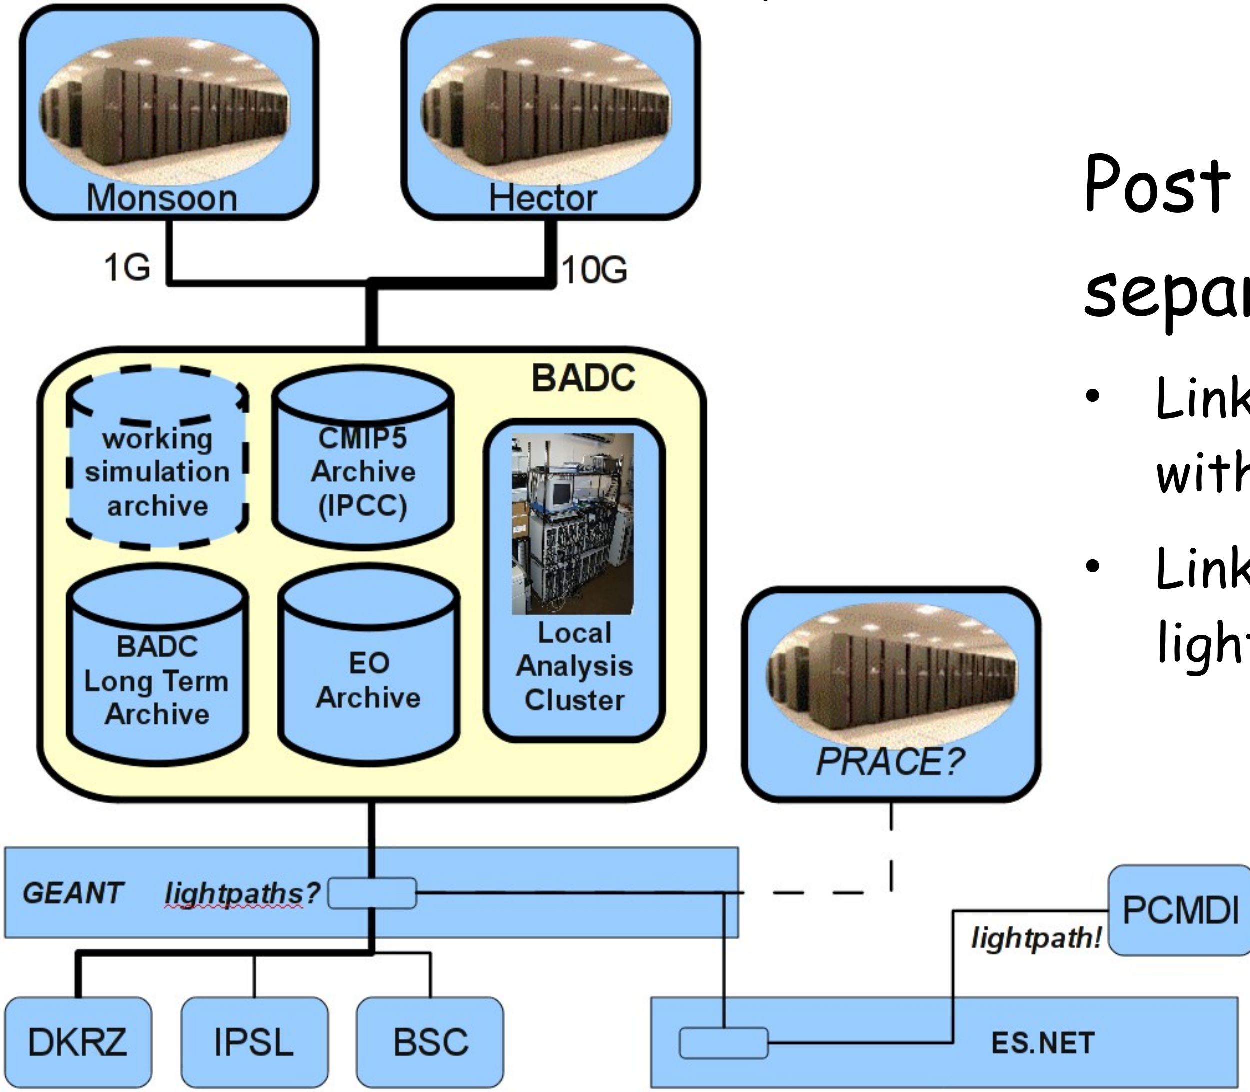

Post processing archive separate from HPC:

- Linked to European archives with lightpaths (if necessary)
- Linked to PCMDI with a lightpath (if possible)

Why light paths?
For CMIP5, synchronising 1 PB archive at 1 % level
- **implies 10 TB/day movement**
- **implies 1 Gbit/s requirement.**

Information matters too: provenance!

At least two reasons:

Practical

- What data should I use for this problem?
  - How was it generated?
  - (For simulations) What parameterisations were included?
- Where is the data?
- How big is it?
- Who was responsible for it? (Reputation matters!)

Evidential

- Who did what, when?
  - Why?
- Can work be repeated? (Similarity is good enough, but similarity requires detail & sometimes input parameters & data)
- I based my conclusions on …
- Who should I cite?

Bringing it all together

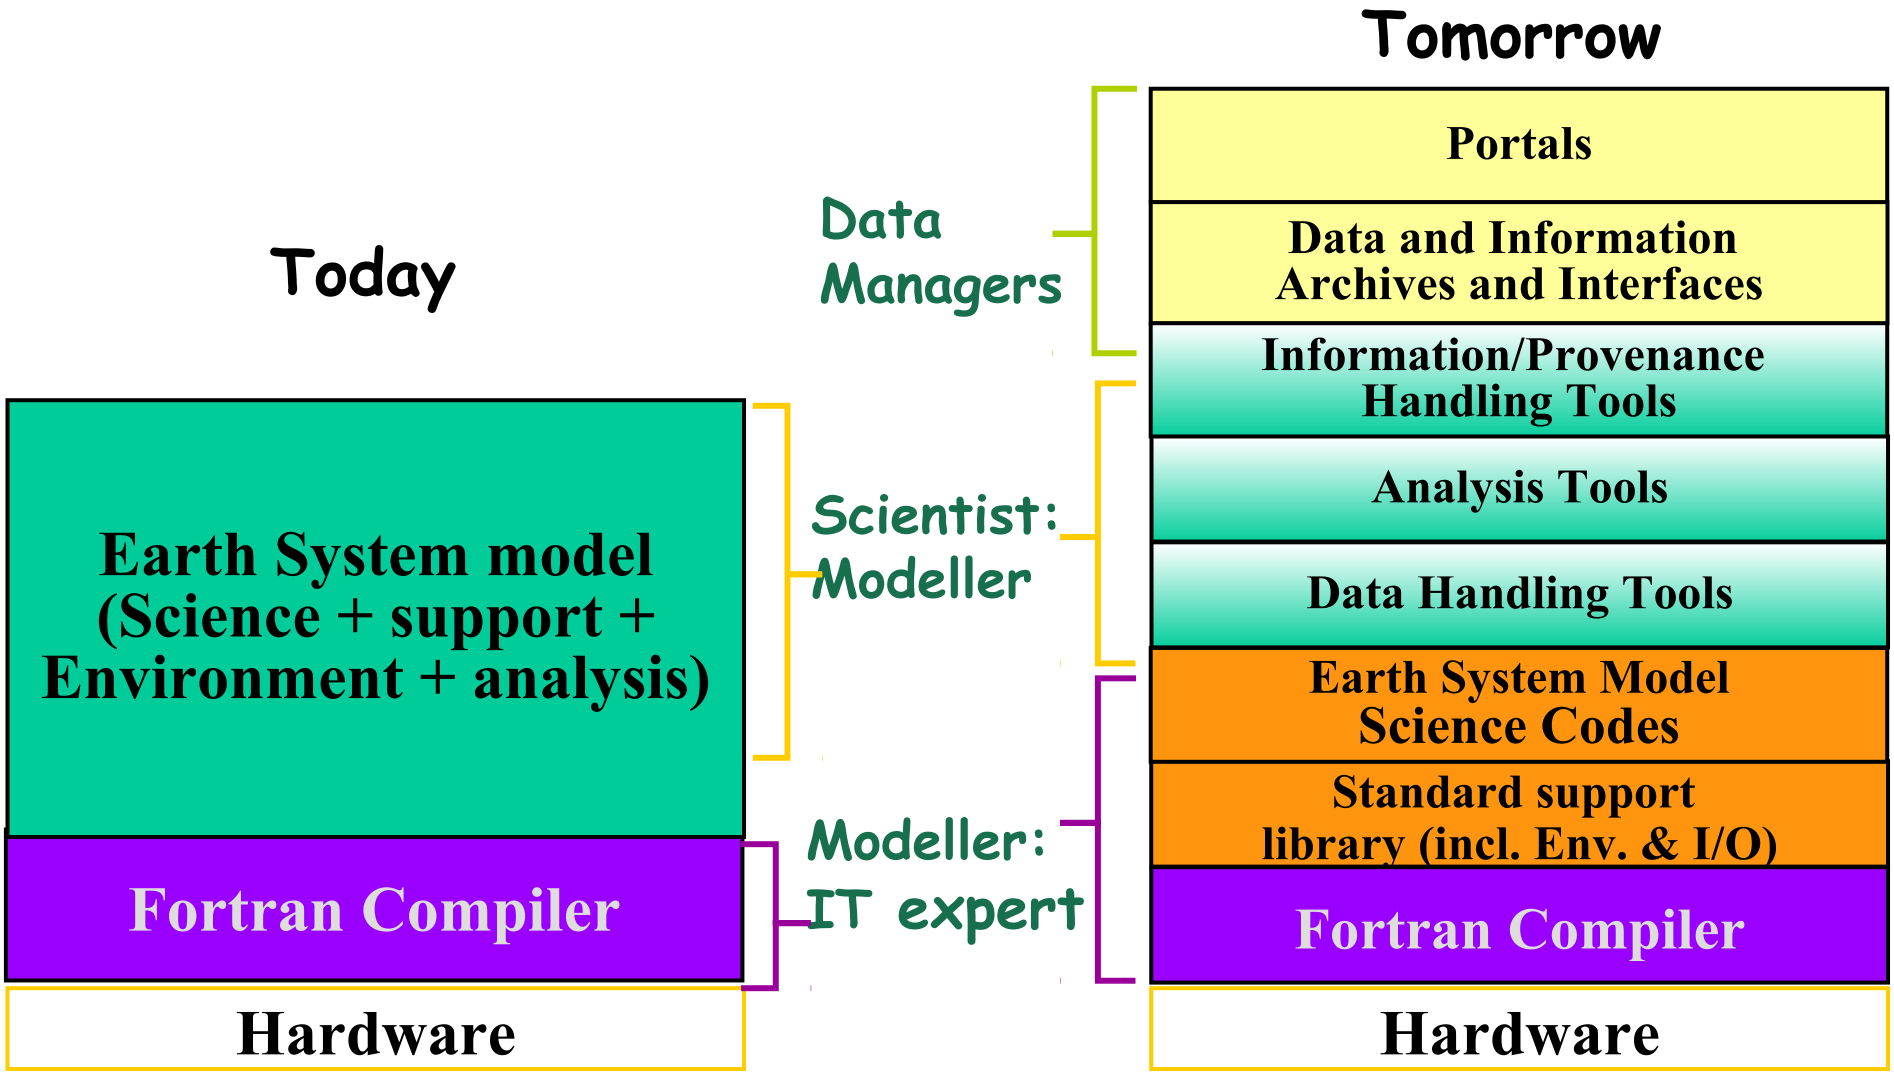

Lots of needs? Some solutions, but?

Solution: Share (much more) ESM software infrastructure across community!

- Share frameworks (support development, maintenance and support),
- Standardise model software environment

Leading to:

- Better performance on a wider variety of platforms, and
- Ease of use of different climate model components.

But is this a pipe dream? Can we govern effectively such a distributed environment?

Wheel reinvention: waste of time or source of innovation?

The science :

- General principles
- Constraints from physical interfaces
- Data Assimilation and Ensemble Analysis
- Input and Output Data

The users:

- Human (GUI) interfaces
- Configuration editor
- Diagnostics outputs
- Tools
- Data Interfaces

Framework

The participating models

- Atmosphere
- Atmos. Chemistry (Aerosol)
- Ocean
- Ocean biogeochemistry
- Sea-ice
- Land-ice
- Land surface

The technical developments:

- System architecture
- Coupler and I/O
- Software management
- Vizualisation and diagnostics
- Automatic Provenance Tools
- Human Provenance Tools

The future

The more we try and work together on ESM science, the more we will need to work together on the software infrastructure within the models, around the models, and associated with the data exploitation and management.

European ESM software needs:

- A network (community) of like-minded people working to common goals (and specifications).
- It already has that (via ENES) with Metafor, IS-ENES and the rump of PRISM group.
- It needs to build on that to exploit what is undoubtedly technical excellence handicapped (still) by geographical distribution and differing local approaches.

Managing the risks

Shared software development needs shared governance backed up a judicious balance of local adherence to shared conventions and local autonomy.

Effective governance

(of conventions, not individuals):

- Incorporates best practice from elsewhere
- Reacts to community feedback and regularises what is effective rather than mandating what is hoped
- Allows evolution (aka innovation)
- Knows when to change direction!

Strategy Recommendations

- Further investigate common modelling frameworks & couplers
- Further investigate common ESM runtime & support infrastructures
- Further support the governance and evolution of key information standards
  - CF and the Metafor CIM and underlying vocabs
- Further support the development and exploitation of key data manipulation tools.
- Look to develop interconnected national archives of large simulations and important observations etc

Do all of this by <u>supporting networks</u> and <u>effective governance</u>

(including recognising the importance of these issues at the ENES steering board level).

The End!

Recap: why common infrastructure?

- European Earth system modelling expertise widely distributed, leads to:
  - Scientific motivation: get more science from facilitating the sharing of scientific expertise and the sharing of models
  - Technical motivation: get more done because the technical challenges are large compared with available effort
  - Efficiency motivation: need to keep scientific diversity (within Europe and globally) while exploiting more complicated software and hardware environments, leads to a requirement for "plug-n-play" (swap in, swap out)!

    (Beware: we geeks know that *unnecessary* paging aka swapping is inefficient!)
- Meeting these challenges is necessary to maintain scientific relevance and competitiveness!
  - Need to provide the best possible policy advice
  - Compare with with US call just last week for 50 million dollars (?annually?) see http://www.energy.gov/news/8777.htm

Managing Data: CMIP5 supporting software

(US) Earth System Grid:

- Data Gateways
  - To provide interfaces to catalogs and show distributed data holdings.
  - Deployed at a handful of locations.
- Data Nodes
  - Expose data for access
  - Deployed with as many of the modelling groups as possible.
- Underlying software
  - ESG Curator
  - Replication
  - Access and authorisation
  - Live Access Server
  - OPeNDAP
  - CMOR2

Global Contributions

- Metafor questionnaire
  - BADC + Metafor team
- Quality Control
  - DKRZ + GO+ESSP team
- Citation Services
  - DOI Allocation (DKRZ)
- More gateways:
  - DOI Portal (BADC+Metafor)
  - IS-ENES
- More underlying software
  - Access and authorisation
  - OGC Interfaces (IS-ENES)

CMIP5, an exemplar: The software runs through it!

The Fifth Coupled Model Intercomparison Project

- Sponsored by the WMO WGCM
- Quality Controlled Data to (eventually) appear in the IPCC Data Distribution Centre
- Data pipeline starts as a WGCM issue, morphs into something loosely ~~governed~~ organised by GO-ESSP, and then become (in part) an IPCC-DDC issue …

• Every European modelling project will/should similar software requirements in the pipeline – from producers to consumers (cf IS-ENES).

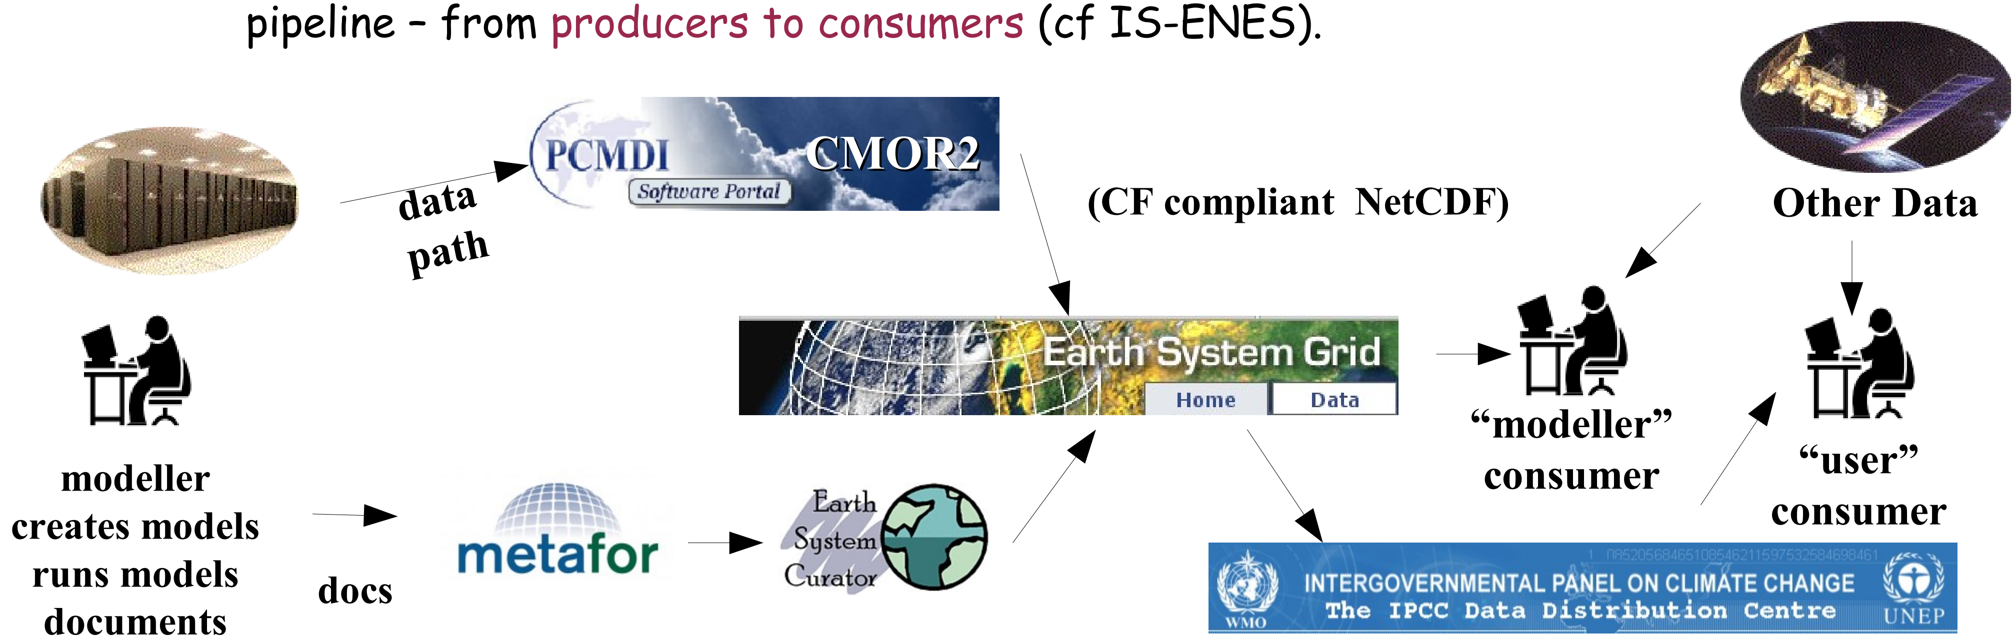

Data processing and visualisation

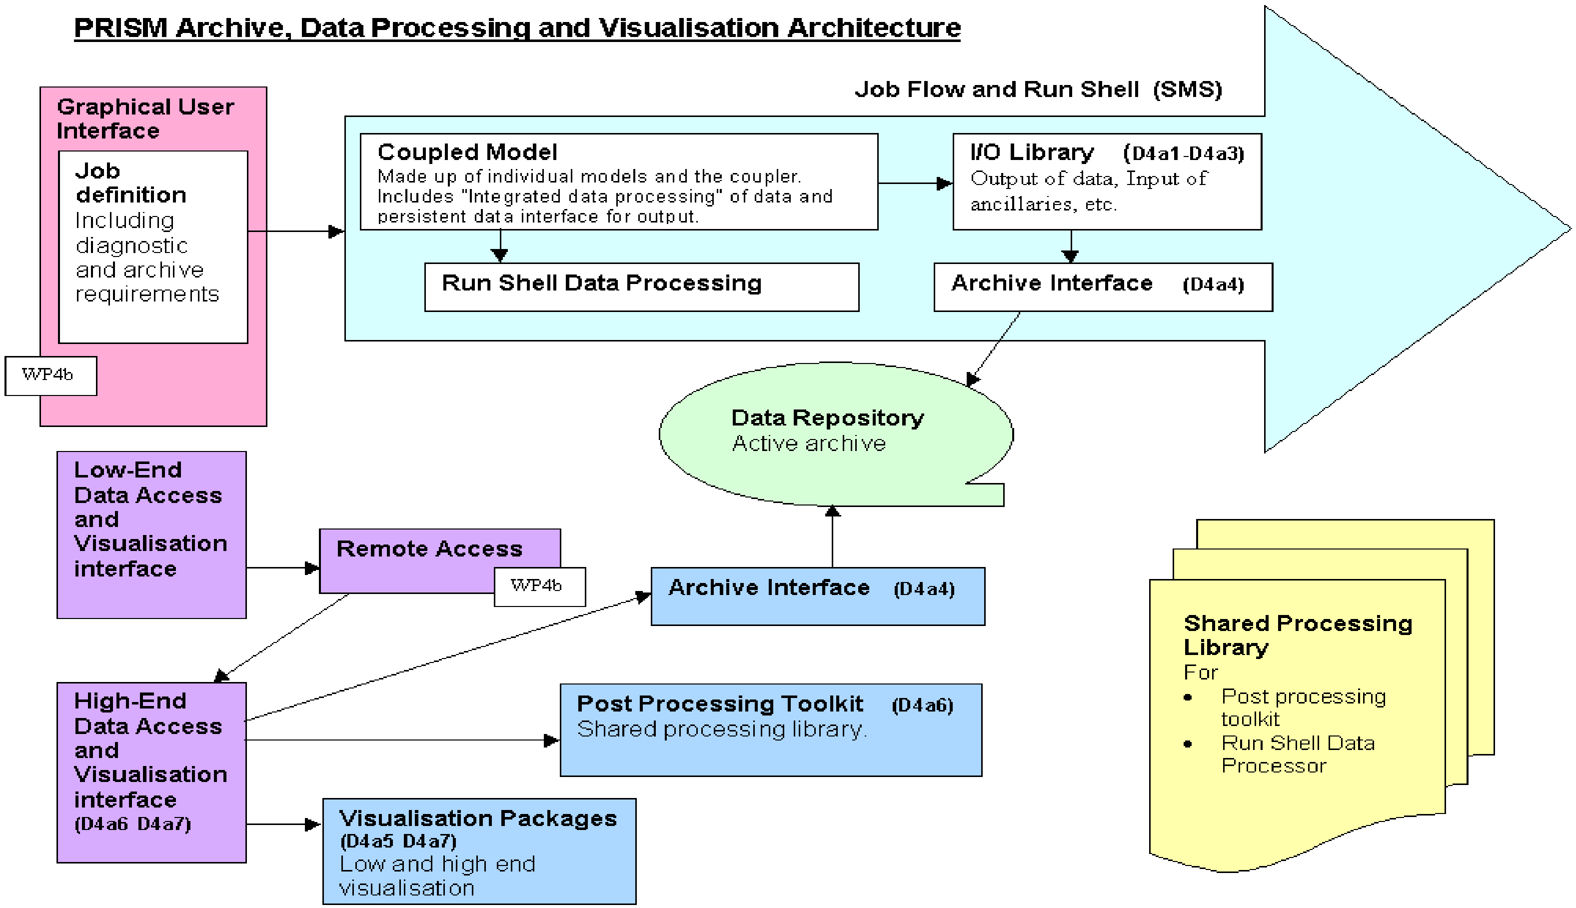

Not necessarily advocating this specific (PRISM) architecture any longer, but we need to something similar!