The Future of ESGF

in the context of ENES Strategy

With a subtext of the important role of IS-ENES2
In addressing solutions to the following question:

"Two thirds of data written is never read!"
WHY NOT?

Bryan Lawrence and friends
(many of whom are in this audience)

Context

Diagram from an early draft of the ENES Infrastructure Strategy (Mitchell et al, 2012):

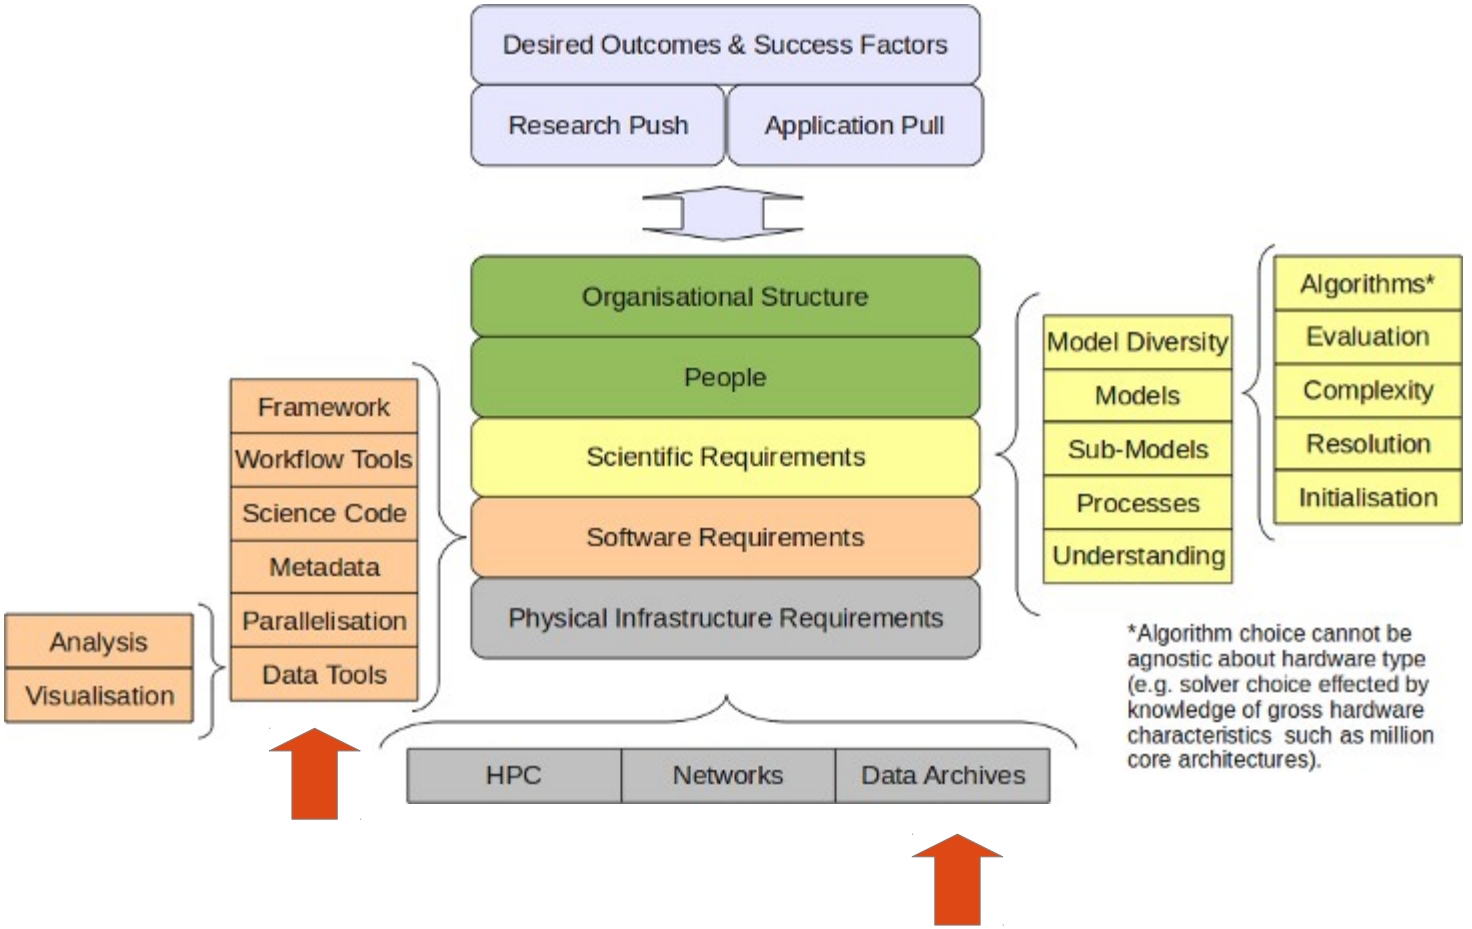

This Talk: Outline

From institutions to Portals, via institutional data archives and ESGF …

- 2/3 of data is never read …
- Status:
  - The consequences of more compute
    - Data at the UKMO and DKRZ
  - ESGF
  - Institutions, ESGF and portals.
- Lessons, Capabilities and Futures
  - Information flow
  - The many roles of IS-ENES2 in the future of the ESGF
  - Data hardware important too!

“Two thirds of data written is never read!”

WHY NOT?

Most simulations are run by individuals or groups only interested in specific phenomena or phenomenon, but not sure what data they need to look at that.

A lot is forgotten!

Potential users aren't sure enough about what the data is, or how it can be used, so they prefer to generate/collect data for which they have provenance!

OR IT'S ONLY READ ONCE, WHY ONLY ONCE?

It is used as input to (one) reductive workflow (e.g. in EO level1-level2, or in simulations, time and/or spatial meaning), then it's being kept “just in case”, and then

As above (forgotten/no-provenance etc), or

It's remembered, but the potential user community no longer have the right tools, or

THEY NEVER HAD THE RIGHT TOOLS! WHAT WAS MISSING?

Hardware

Software

Information

- (did they know that data of interest to them even existed?)

Do we expect anything different from our portals?

Bryan's corollary:

2/3 of portals will never be used.

unless they attend to **user requirements**
and provide **adequate quality control**
and **provenance**.

(using the well known fact that 2/3 of all statistics are made up on the fly)

Solutions

1) Curation (of data)
   - So it wont be forgotten and tooling or data (or both) are migrated with time.
2) We don't care:
   - We can rerun simulations (or, in rare cases, re-collect data).
3) Build better tools.

All options depend on

Metadata:

- In the case of **curation**, so the data can be curated, and migrated, and users can understand the provenance etc.
- In the case of **re-running,** so the same data can be re-produced (so we need descriptions of the simulator, the inputs etc, and if we're really picky, we need the same computer and compiler!)
- In the case of tools, to support reusable and **parallelisable workflows.**

Resources:

- **People, Software, and Hardware** (in that order)!
- AND:
- People doing “the right thing”.

But note that none of those address
the portal problem directly!
All necessary, but
NOT sufficient!

And the problem is only going to get worse!

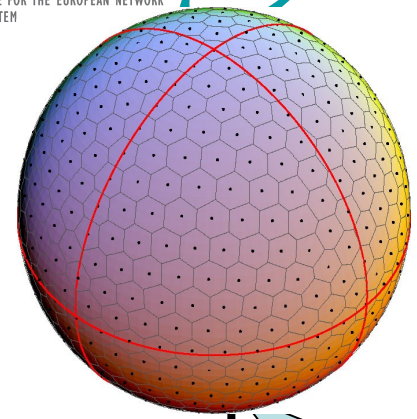

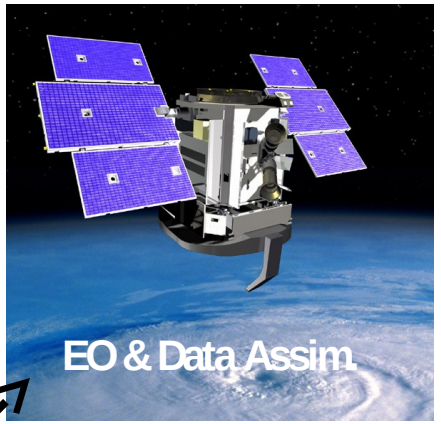

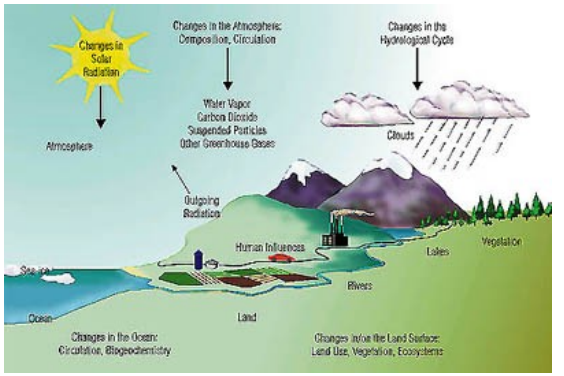

Resolution

Enhanced computing resources produce MORE DATA

Complexity

Duration and/or Ensemble size

(Many versions of this slide exist, this one from J. Kinter's presentation to the world modelling summit 2008)

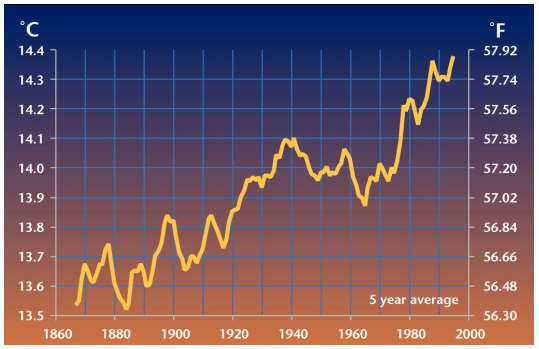

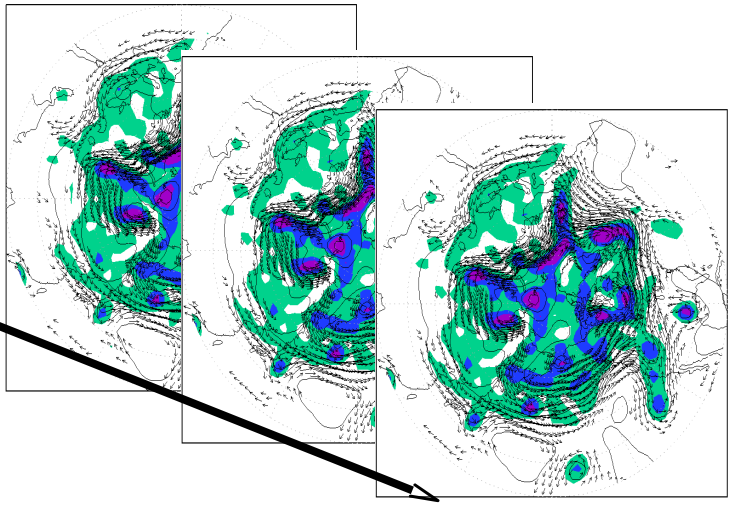

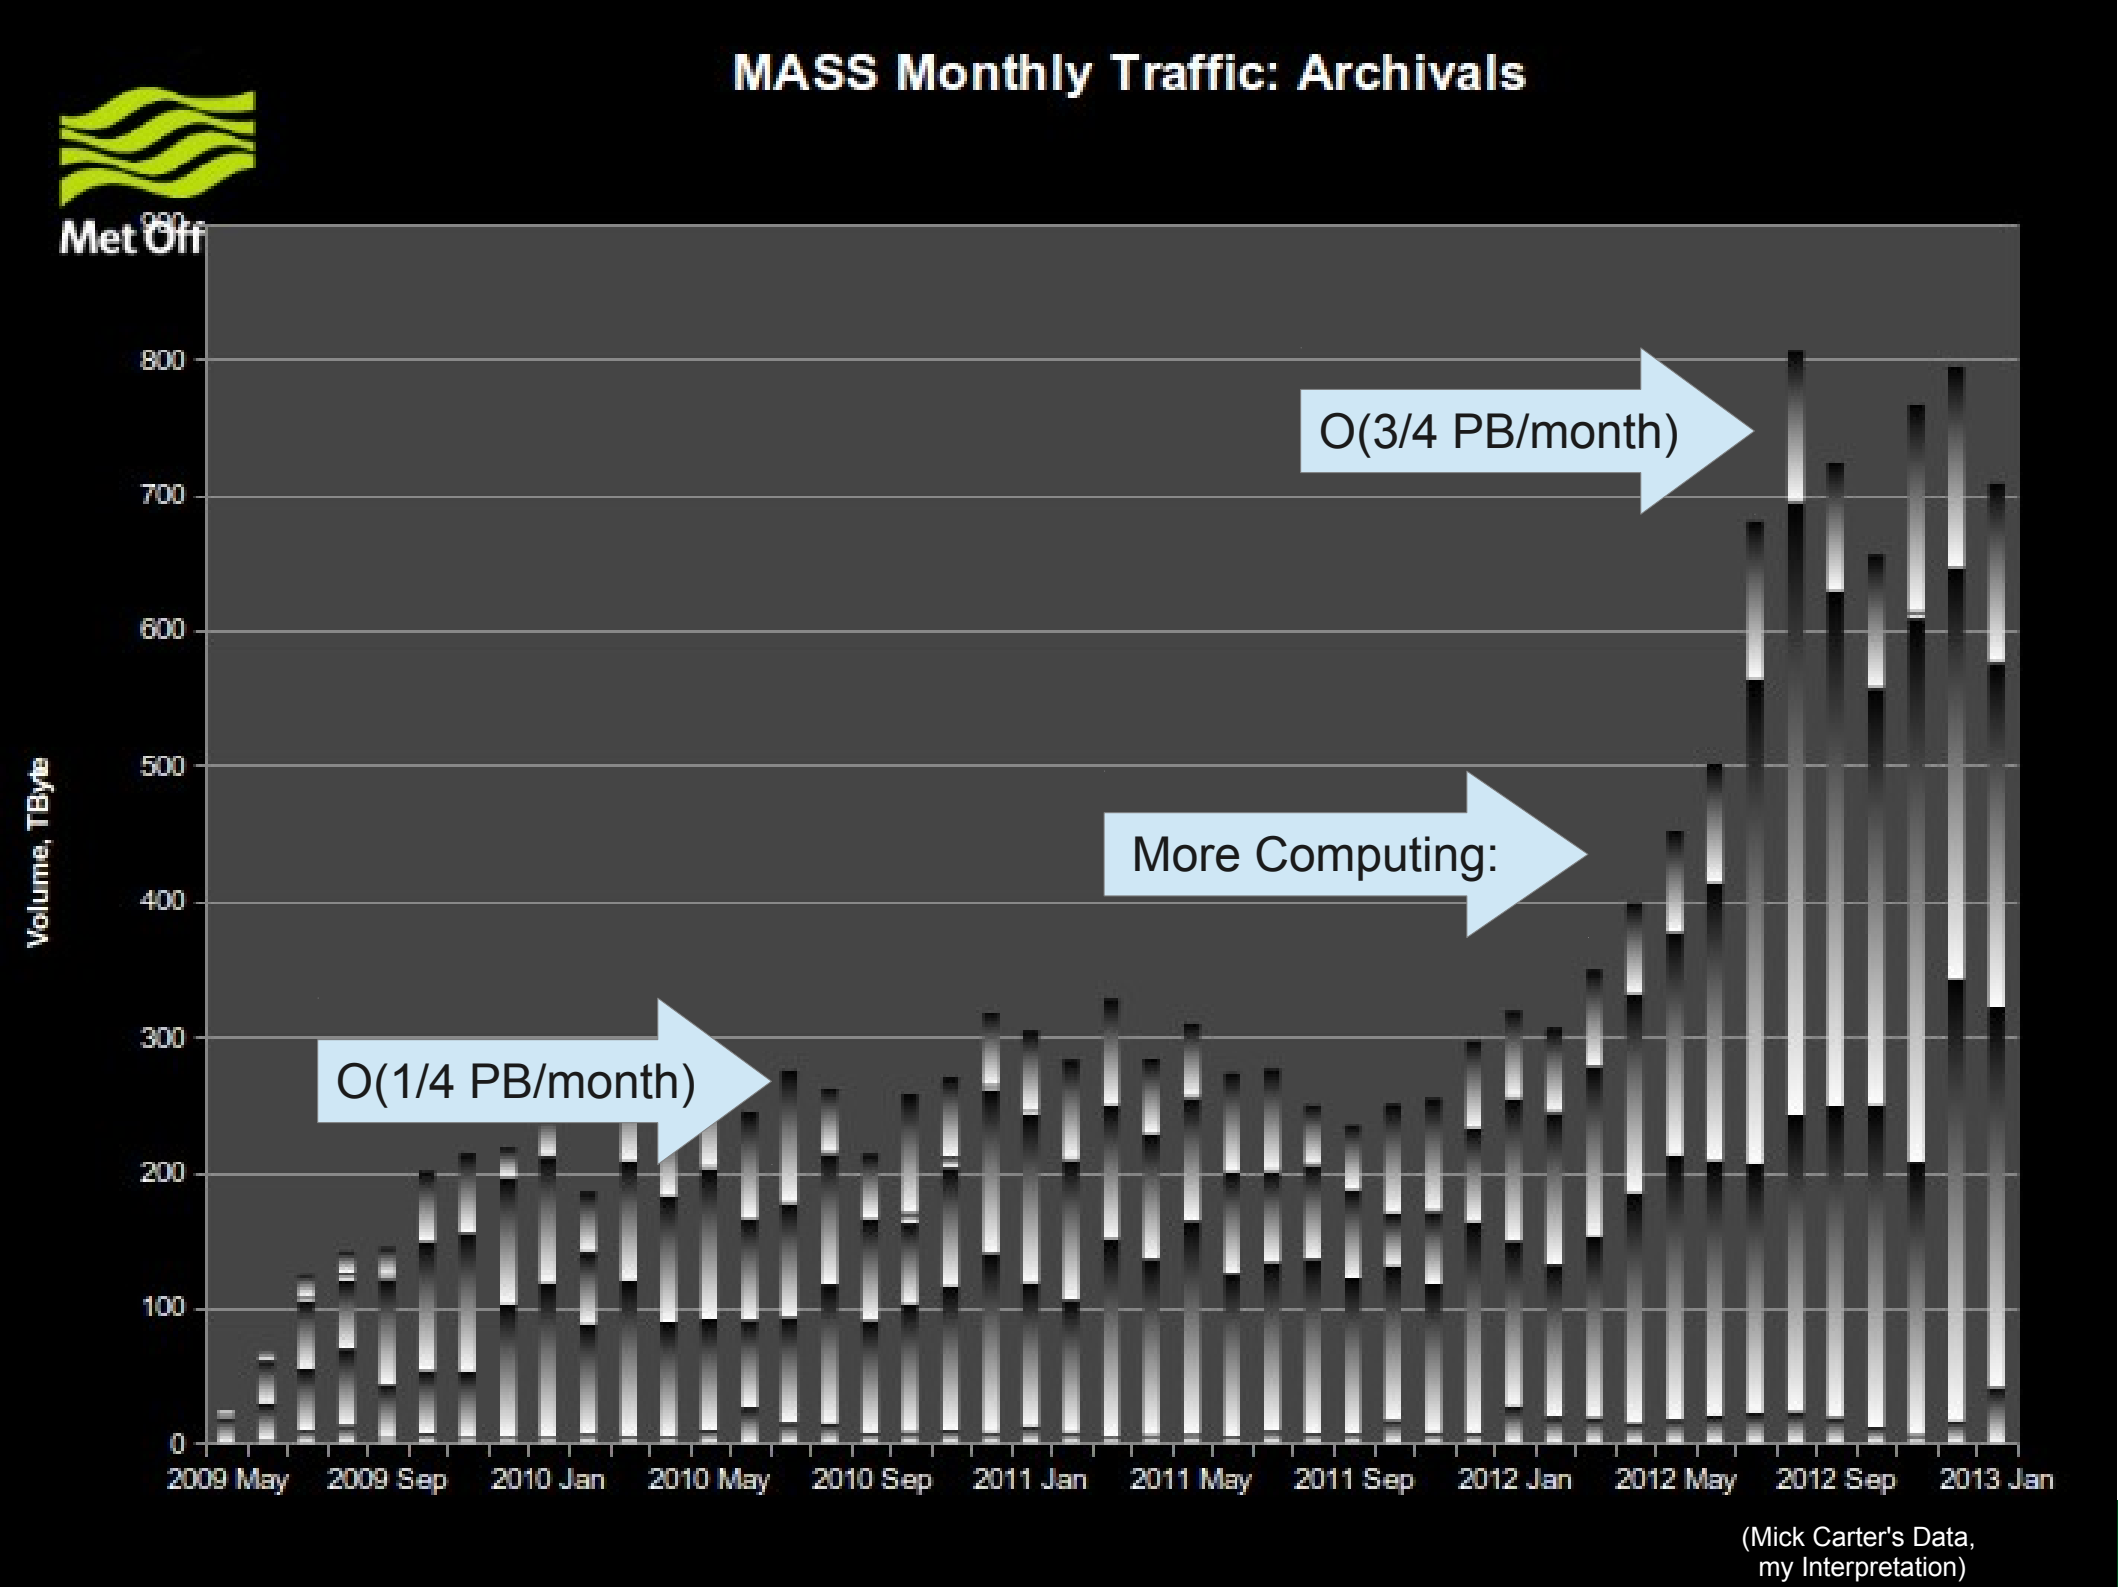
MASS Monthly Traffic: Archivals
Met Office
Volume, TByte
800
700
600
500
400
300
200
100
0
2009 May
2009 Sep
2010 Jan
2010 May
2010 Sep
2011 Jan
2011 May
2011 Sep
2012 Jan
2012 May
2012 Sep
2013 Jan
O(3/4 PB/month)
More Computing:
O(1/4 PB/month)
(Mick Carter's Data,
my Interpretation)

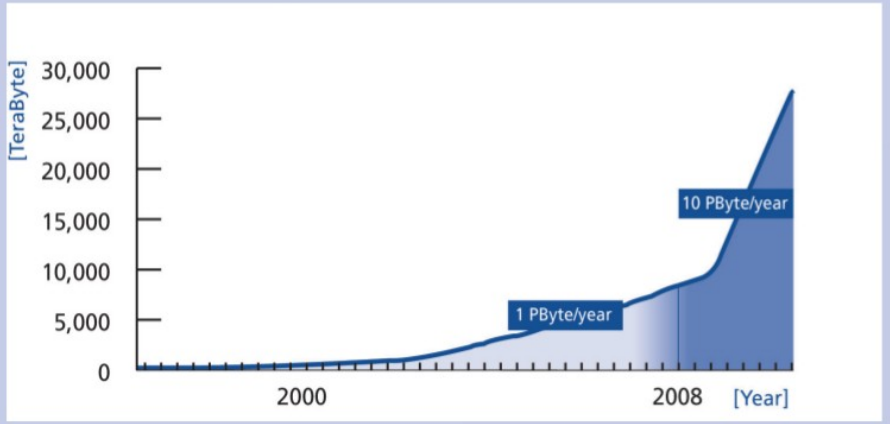

Exponential data growth at DKRZ

The DKRZ archive at a glance

- 7 automatic Sun StorageTek SL8500 tape libraries
- 8 robots per library
- more than 65,000 magnetic cartridges
- 78 tape drives
- total capacity: over 65 PetaByte
- bandwidth 5 GigaByte/s (bidirectional)

Numbers

Currently DKRZ produces 10 PetaByte, i.e. 10,000 TeraByte of data per year. This corresponds to two million video-DVDs. The data can be written and read with a bandwidth of 5 GigaByte/s, i.e. every second the content of one DVD can be transferred from disk to tape.

(From DKRZ poster at ISC 2011, courtesy Michael Lautenschlager)

We don't have much of a history of curbing our appetite for writing data!

We do have a history of putting lots of effort into our mass storage architectures at the “big centres”, and we're now moving onto federated data systems.

Big Data Locally and Globally

The Earth System Grid Federation

Data Nodes,
providing data services
and publishing to
Data Indexes/Gateways
linked in a
Global Federation

At least three nodes committed to ***persisting*** the data!

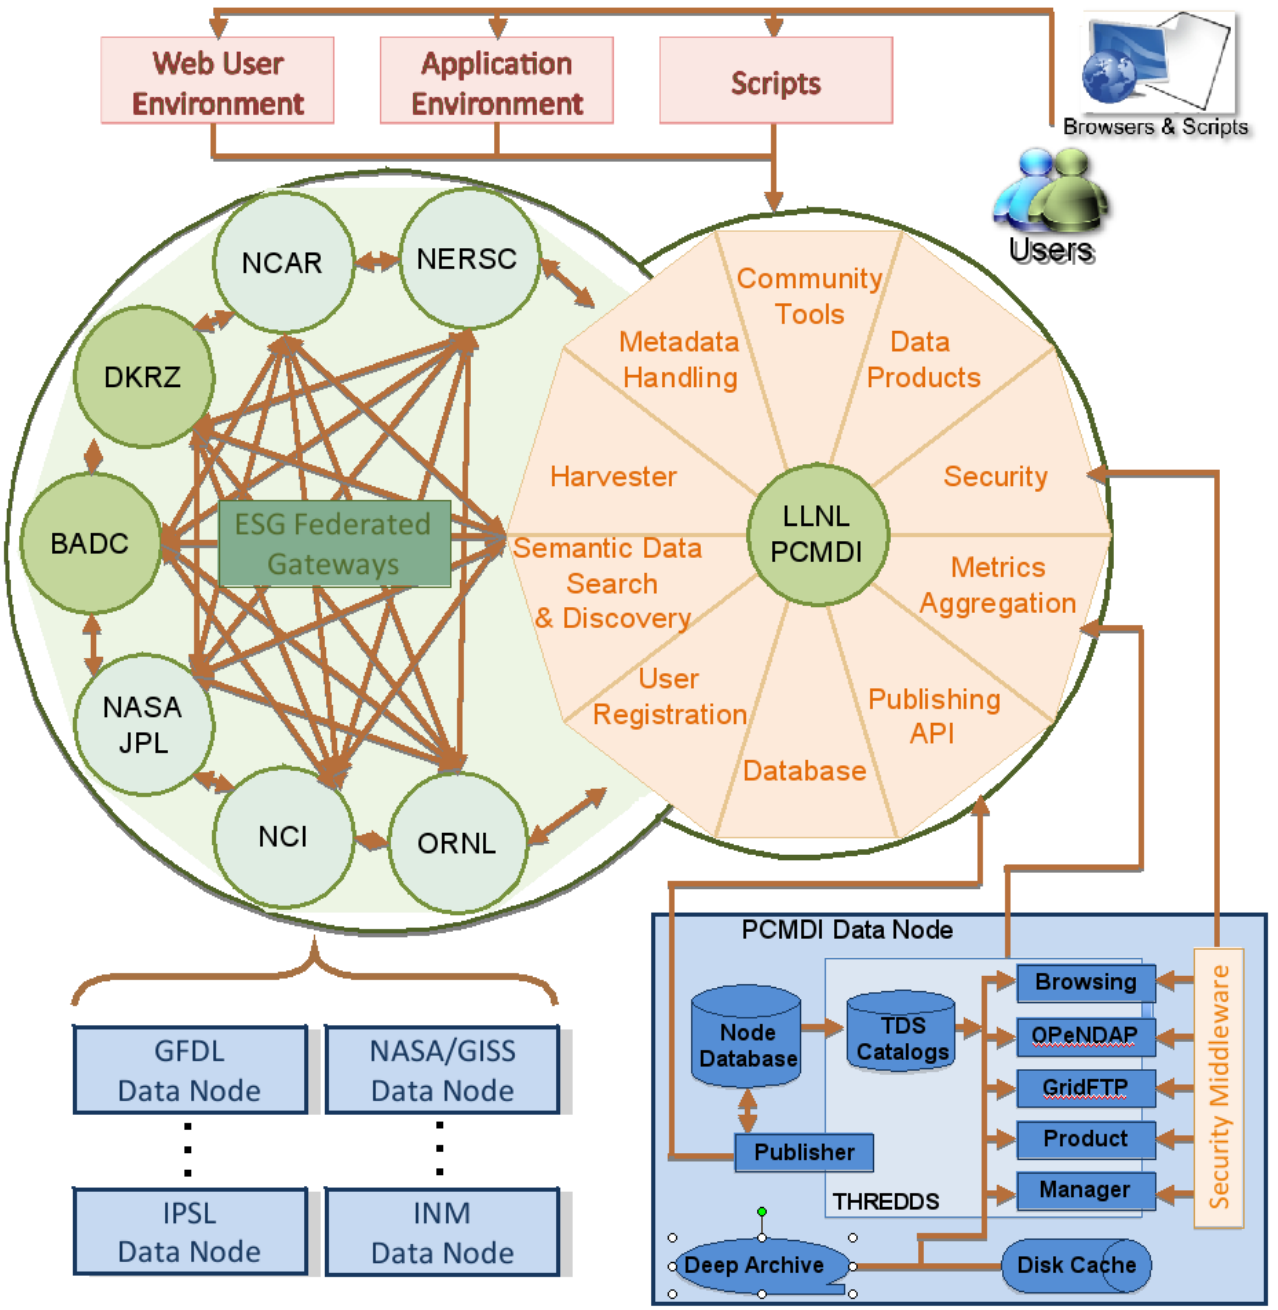

is-enes
INFRASTRUCTURE FOR THE EUROPEAN NETWORK FOR EARTH SYSTEM MODELLING

enes European Network for Earth System Modeling
THE ENES PORTAL
Service for Climate Modeling in Europe

CAPACITIES

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

DKRZ

CMIP5 European index gateways
Access
Quality control
Replicas

CIM

metafor

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

CMCC
Centro Euro-Mediterraneo per i Cambiamenti Climatici

DKRZ

ICHEC
Irish Centre for High-End Computing

Institut Pierre Simon Laplace

Norwegian Climate Center

CNRM-CERFACS

METEO FRA
Toujours un temps d'a

CMIP5 European data nodes (but there are more)

PCMDI

ESGF
Earth System Grid federation

ENES collaboration with PCMDI, ESGF and GO-ESSP

Where next with ESGF?

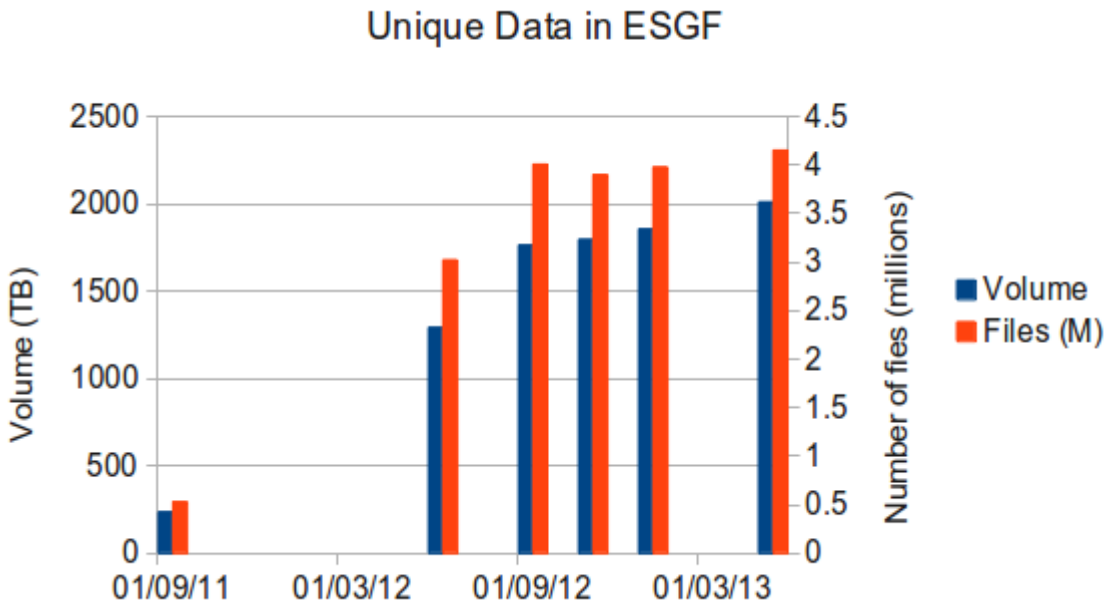

CMIP5 Archive Status

Last Update: *Thursday, 23. May 2013 12:11AM* (UTC)

CMIP5 Federated Archive

Summary	
Modeling centers	29
Models	61
Experiments	101
Data nodes	23
P2P Index	11
Datasets	59082
Size	2,008.37 TB
Files	4,150,032

Latest version only; no replicas.

Lots of new activities joining in, but unlikely to be very large deltas until CMIP6 (except, possibly, adding significant amounts of UPSCALE data.)

ESGF Dependencies

ESGF depends on

- Constraining data provision to specific formats (NetCDF), and specific conventions (CF+CMIP5 specific constraints).
- Metadata conventions.
- Constraining the data layout on disk. A filename convention.
- Agreements on how to do, and use, authentication and authorisation (openID+X509).
- A lot of opensource software!

ESGF: Missing Pieces

Model Documentation

CMIP5: 2 PB, 101 experiments, 61 models.

- Crucial role of model documentation, now being fully integrated into ESGF based on Metafor and IS-ENES.
- But much yet to do, better tools for creation and reporting, significant role for IS-ENES2!

Quality Control

CMIP5: 60 thousand datasets and 4 million files.

- Impossible for users to find and correct **all** mistakes, or deal with format inconsistencies without excessive effort.
- Caveat Emptor is not enough, significant burden on data providers is necessary!
- Reward that provider burden with “data publication”.

Service Portals

ESGF provides data download, with some (mostly unused) Live Access Server capability.

- It's designed for climate modellers!
- Most data nodes don't have the horsepower to deliver “live services”.
- Major role for IS-ENES2 in developing and deploying service-orientated portals for real users, but it will be important to ensure they have the capability to deliver!
  - The right information to the right people via the right interfaces

Bringing all those things together

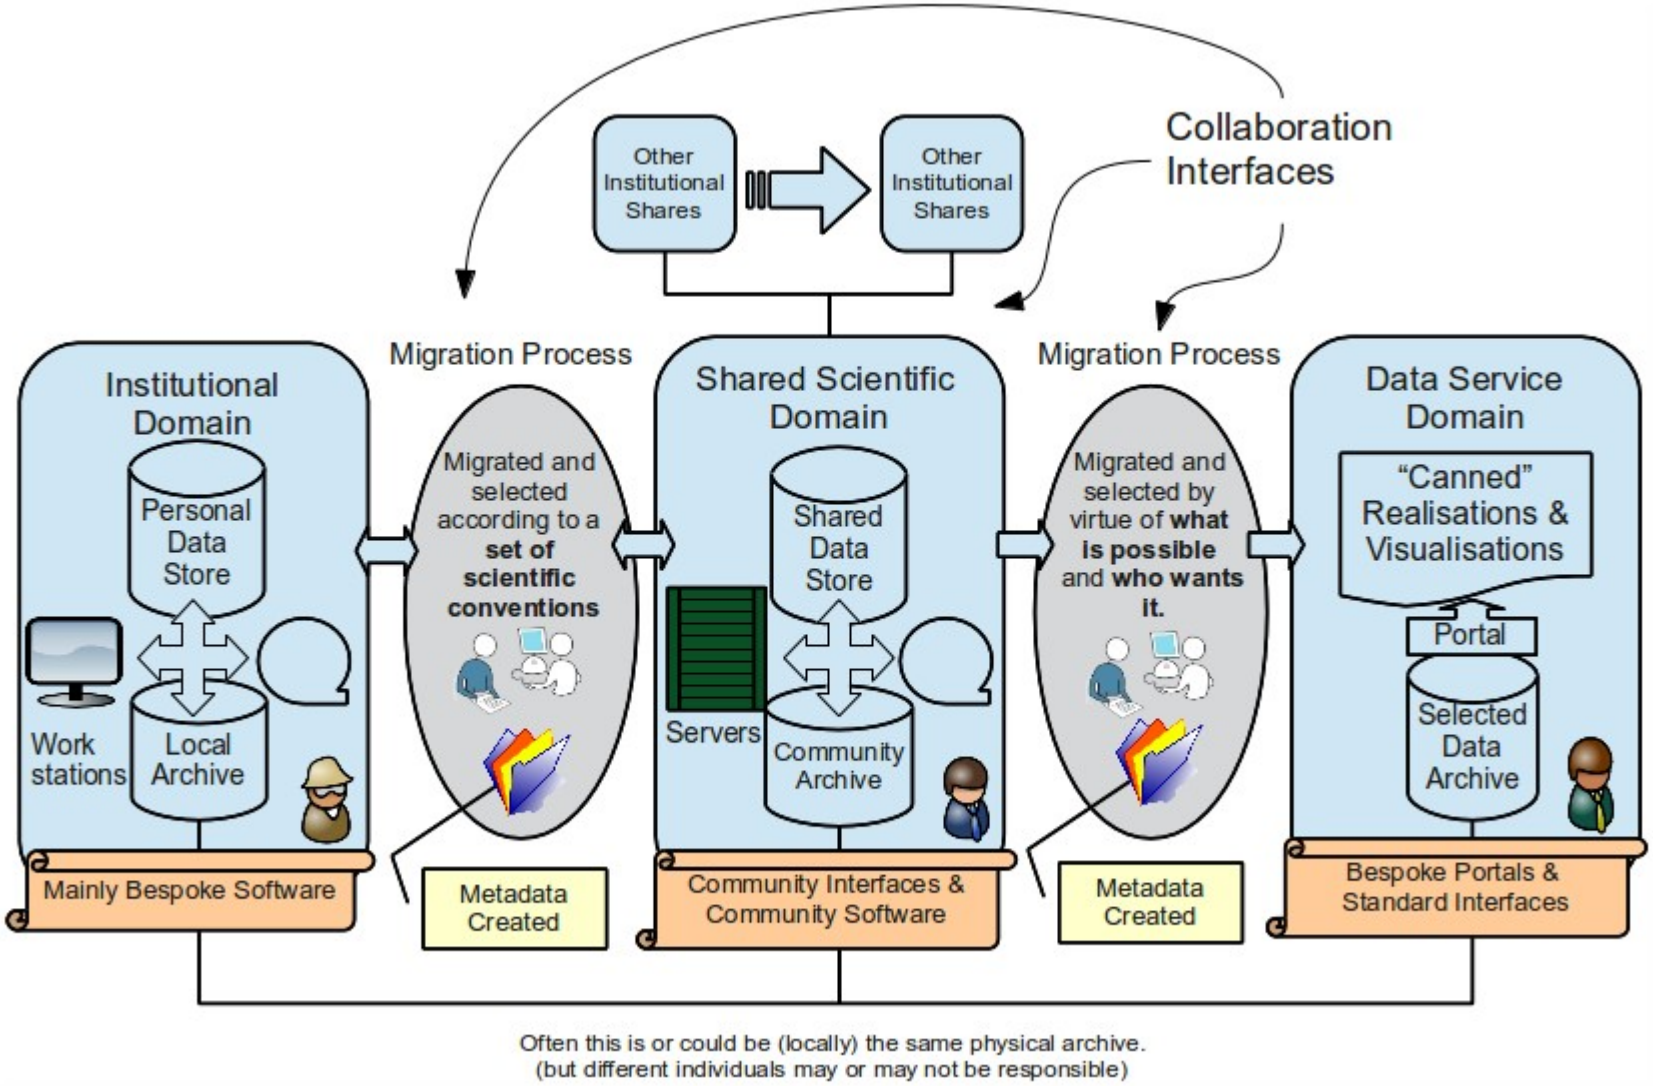

From concepts in Treloar et al 2007 (D-LIB)

IS-ENES2 Version

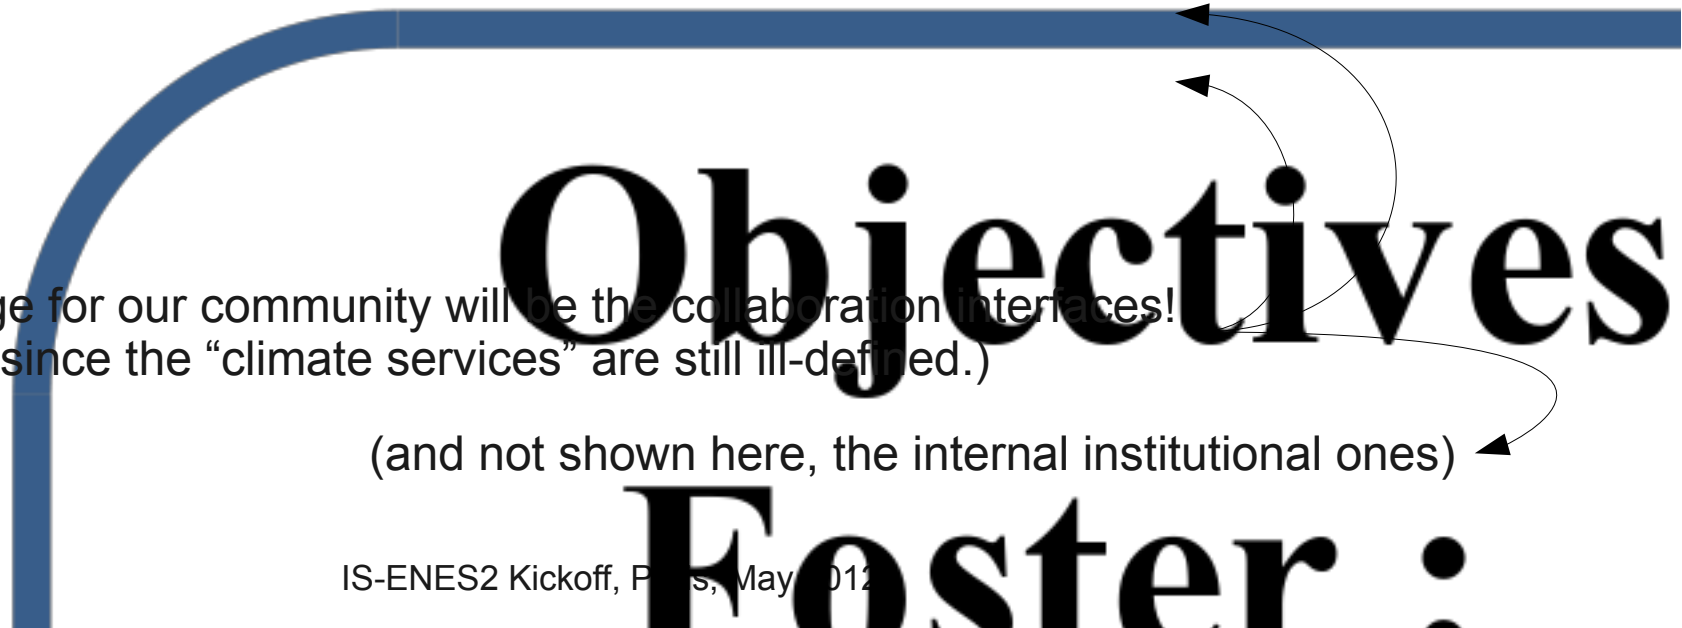

The challenge for our community will be the collaboration interfaces!
(Particularly since the “climate services” are still ill-defined.)

(and not shown here, the internal institutional ones)

Data Infrastructure Depends on Information Infrastructure

We're dealing with data for and in three domains:

- Institutional (Science)
- Federated (Science, but different community)
- Shared (Services)

Between each we have collaborations, formal or informal, and between each, we have more and more layers of information, conventions and common behaviours!

The hard part is that the producer on the inside doesn't always know (or want to produce) what the next layer up wants or needs.

Crucial to remember that in this environment “if it isn't in the metadata” the outer layers “WONT KNOW IT”!

- Can't rely on institutional wisdom and/or notes on portals which don't flow outwards!

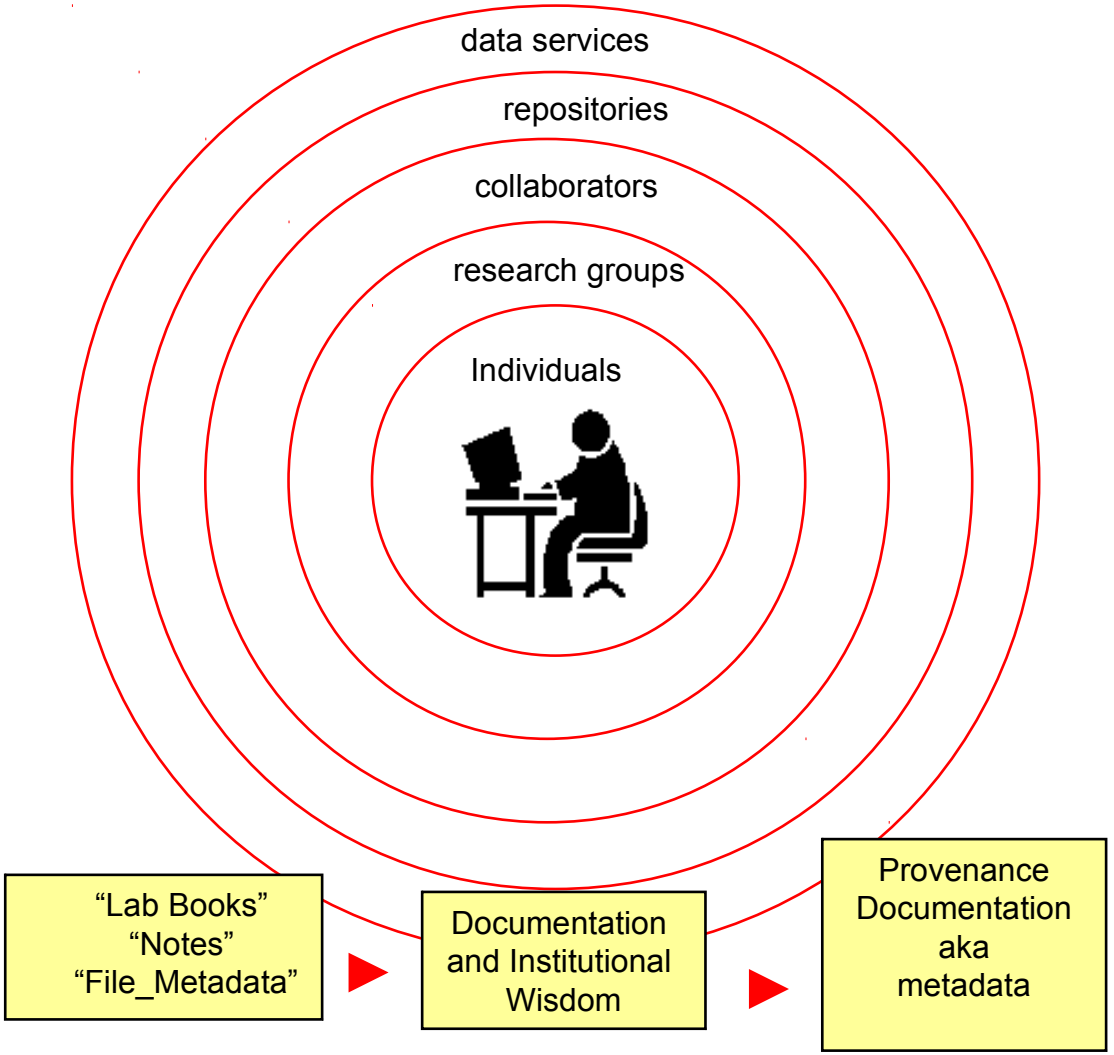

Lessons from History

Every CMIP5 phase has increased the emphasis on consistency of data formats, and quality of metadata. Why?

- Because doing so is crucial to exploiting the data

Every CMIP modelling site has struggled to format and document their data. Why?

- Because: (generally) the activity is under-resourced and/or undervalued!

Conclusion: Data handling for community collaboration is becoming a larger, and more important component of the community workflow.

The ESGF *organisation* has been delivered by a coalition of the willing, using software developed by those *who could.* The entire edifice is powered by soft money.

- No one thinks this is a good idea (viz the US national academy report on climate modelling, and our own ENES foresight.)
- There are governance issues: both for the organisation, and the software!

Conclusion: Significant role for ENES in ongoing activities and governance!

ESGF and new communities

ESA-CCI

- ESA Climate Change Initiative
- European equivalent of OBS4MIP.
- Much effort into developing appropriate formats, and compliance.
- Will join ESGF under the auspices of (other) EC funding.

CORDEX

- COordinated Regional climate Downscaling Experiment
- Needs to work on format conventions and compliance.
- Nearly in ESGF now, will be supported by IS-ENES2.

Expect a spectrum of candidates to put data into ESGF;
It will be important to hold a line on format compliance and quality control!
- but who holds that line?
IS-ENES2!

ESGF & New Software

Significant role for ENES in requirements capture and development!

ESGF software has *nearly* reached a stage where there are defined APIs!

It will be crucial for ENES for it to *actually* reach such a stage, so we can evolve ESGF according to *our* requirements.

We can hope that the current uniformity of the ESGF portals will be replaced by a functional diversity, better targeted to our user communities.

Governance of the software will have to operate on a different level to the governance of the federation.

Important roles for ENES, e.g. JRA3:

- improve version control and package management.
- improve metadata tooling, including for QC and annotation.
- monitoring systems!

Data Services

The US National Academy of Science's “National Strategy for Advancing Climate Modeling” (2012):

The United States should develop a national information technology (IT) infrastructure for Earth System climate observations and model data that builds from existing efforts, so as to facilitate and accelerate data display, visualization, and analysis both for experts and the broader user community.

(BNL: An infrastructure that supports all of these things doesn't have to do all these things in one portal or software stack!)

Without substantial research effort into new methods of storage, data dissemination, data semantics, and visualization, all aimed at bringing analysis and computation to the data, ***rather than trying to download the data*** and perform analysis locally, it is likely that the data might become frustratingly inaccessible to users.

So what would that infrastructure look like?

It would have software tools for massively parallel analysis of data, and those tools would support both:

- Parallelisation within institutional boundaries
- Parallelisation by distributing tasks across (the new) ESGF,and
- Parallelisation by caching large quantities of data next to massively parallel computational resources, including behind portals.

Projects working on these concepts include

- the G8 ExARCH project,
- future incarnations of ESGF being developed at PCMDI,
- IS-ENES2 itself,
- National activities, such as (in the UK) JASMIN/CEMS,
- Institutional activities.

But nearly all suffer from being “soft” money software developments, which mean they have to keep evolving. Only some national activities have “recurrent operational” support – ***hence important role for IS-ENES2 services!***

Next Generation JASMIN

Upgrades planned
2013/14
and
2014/15

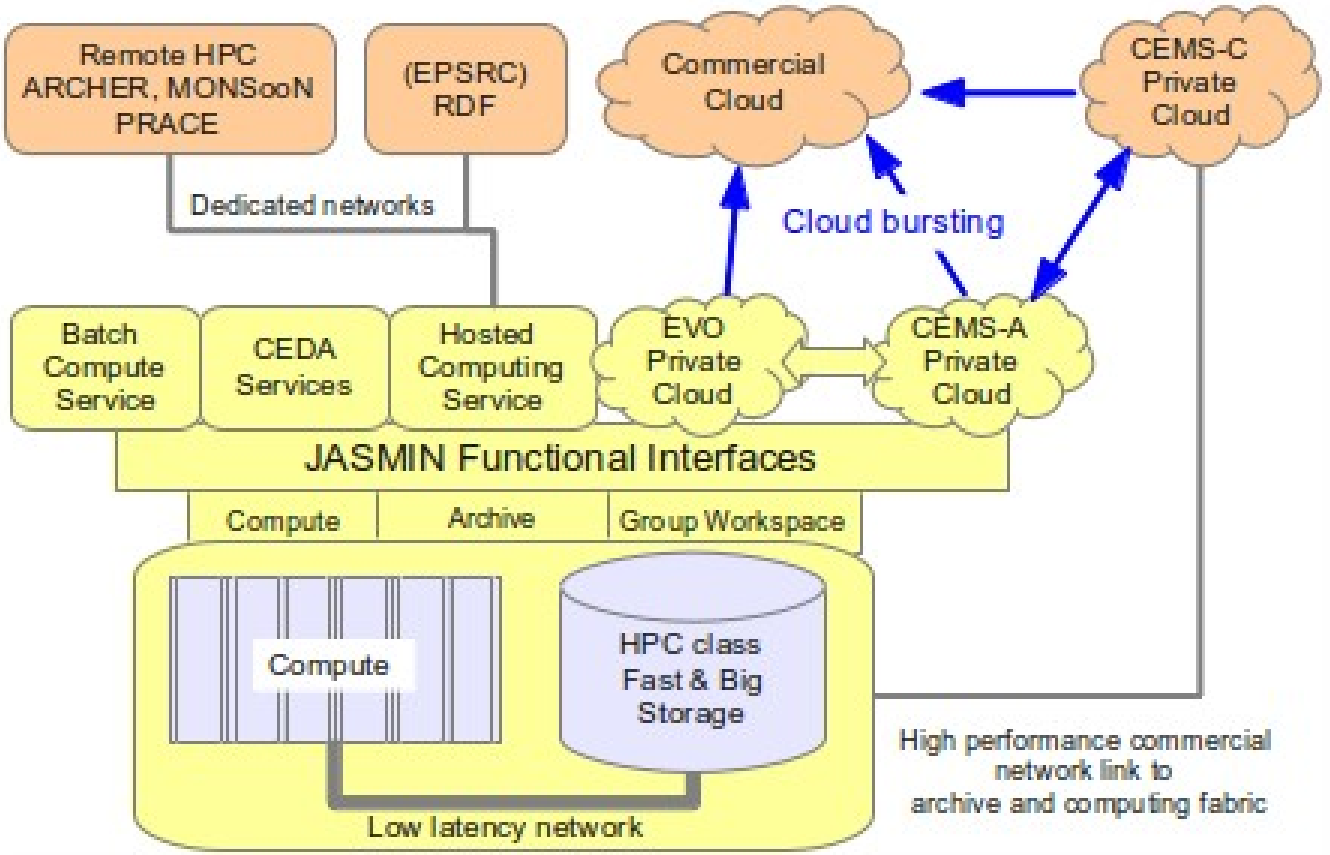

~11 PB spinning disk, plus elastic tape (tens of PB)
2000 cores + support to cloud burst to Amazon/Microsoft et al
Designed to support all three domains:
institutional science, shared science, and (climate) services

Summary

Data Handling is a crucial part of the climate modelling ecosystem! And it depends on:

- Collaboration!
- People (to do migrate data and create information)
- Software (which needs investment and governance)
- Organisation (which needs governance)
- Community (buy-in: both in terms of usage and doing the extra work in data documentation and quality control).

ESGF provides part of that ecosystem, but we need to recognise some limitations (both currently, and in general – it's not the ecosystem in it's entirety).

The future of the ESGF?

- The “G” depends very much on IS-ENES2!

MORE CONTEXT

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?