Coopérations internationales, simulations climatiques et masse d'informations

Sous titres:
Les “challenges” en termes d'accessibilité et d'usage

Sébastien Denvil (IPSL/CNRS)

Outline

- Motivation
  - Drowning in data, with more to come.
- CMIP5 as an example of the current technology
  - Different views; institutional, federated, subsetted
  - Have we hit the limit for downloading?
  - Metadata infrastructure
- The future
  - Taking the computing to the data
  - Taking the computing and the data to the data
  - Scaling to the future requires standards and conformance

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

prodiguer

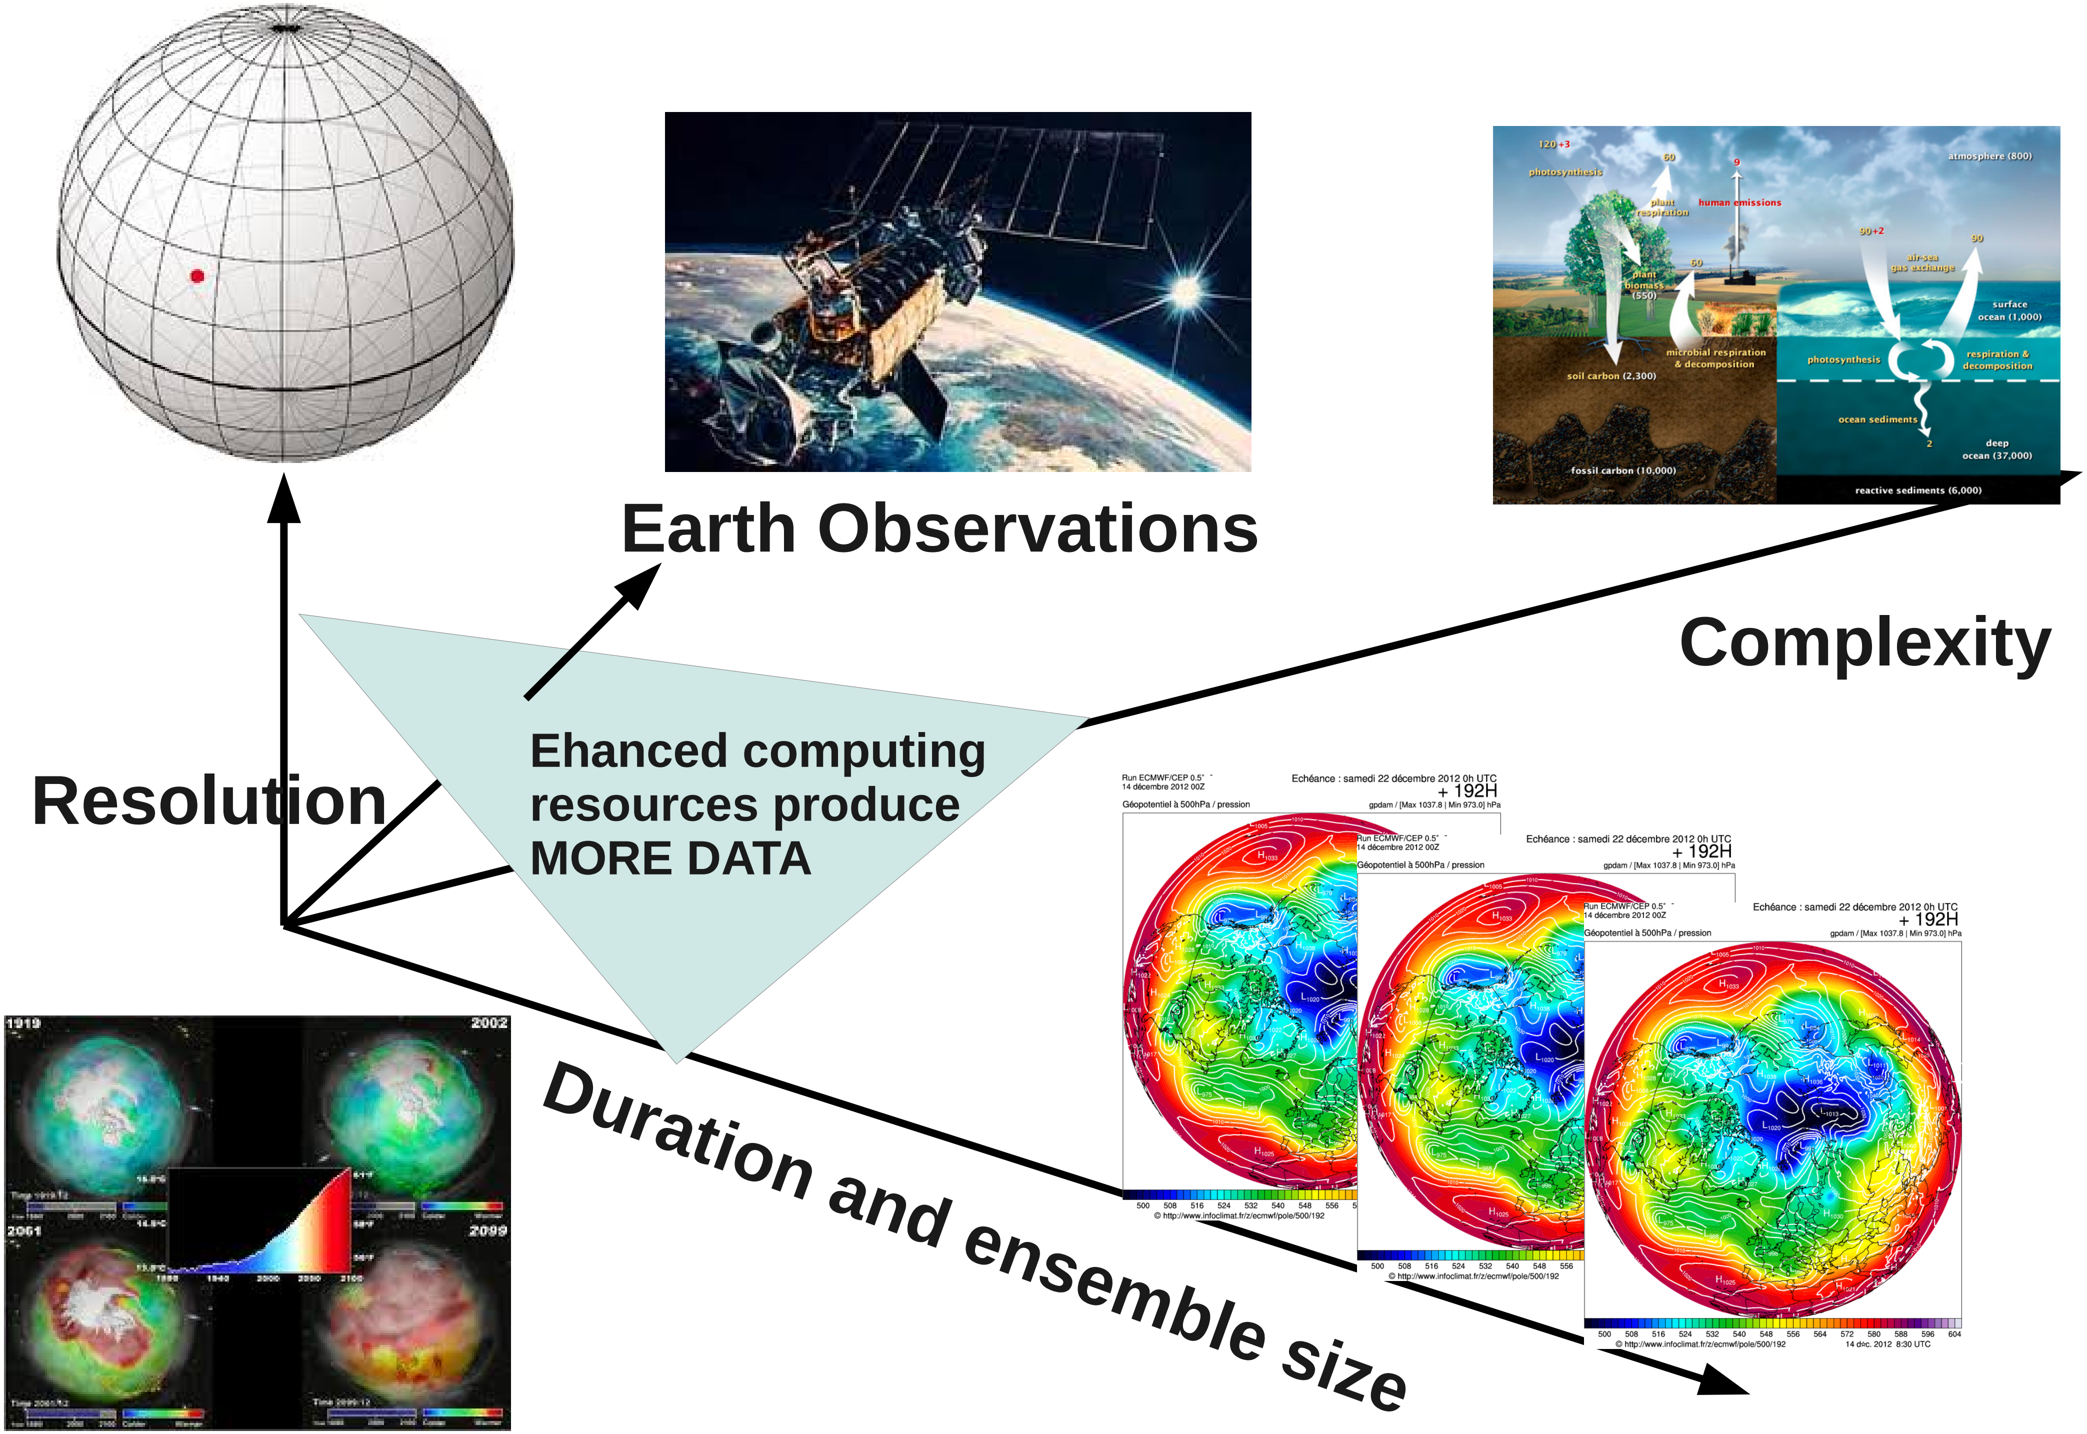
Resolution
Earth Observations
Complexity
Ehanced computing
resources produce
MORE DATA
Duration and ensemble size

To keep in mind

“the potential to interpret, compare and reuse climate information results is strongly related to the quality of their description”

Computation useless if results cannot be stored/distributed/read

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

prodiguer

Many, many processes, many, many communities !

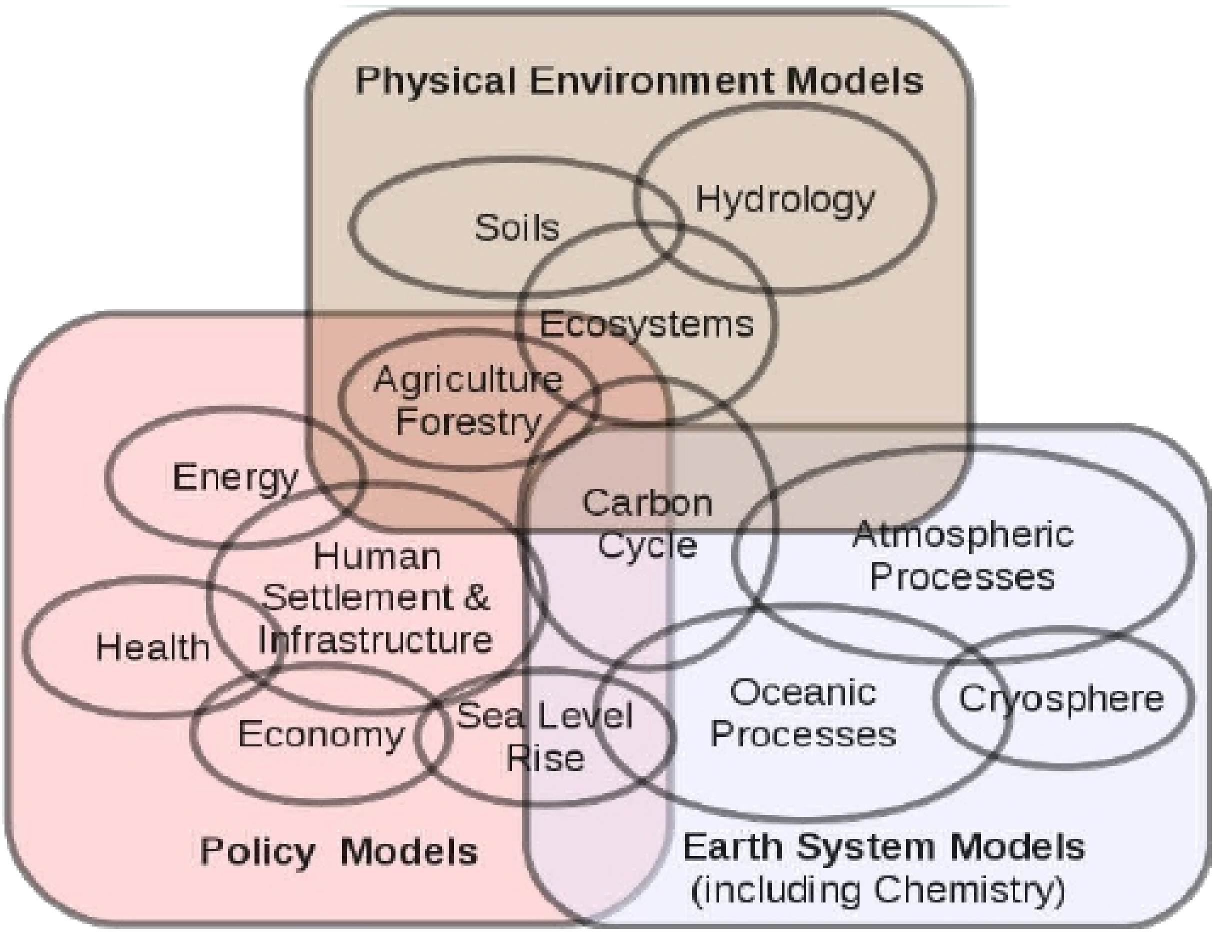

Interconnected communities, all needing access to (some) of the data!

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

prodiguer

IPCC AR5 variable counts

	1 hour	3 hour	6 hour	daily	month	annual	**totals**
aerosol	0	0	0	0	81	0	**81**
atmosphere	75	101	9	86	184	0	**455**
land	0	3	0	2	59	0	**64**
land ice	0	0	0	2	13	0	**15**
ocean	0	1	0	3	116	0	**120**
biogeochemistry	0	0	0	0	88	71	**159**
sea ice	0	0	0	4	47	0	**51**
totals	**75**	**105**	**9**	**97**	**588**	**71**	**945**

Overpeck et al. Science 2011;331:700-702

Fig. 2 The volume of worldwide climate data is expanding rapidly, creating challenges for both physical archiving and sharing, as well as for ease of access and finding what's needed, particularly if you are not a climate scientist.

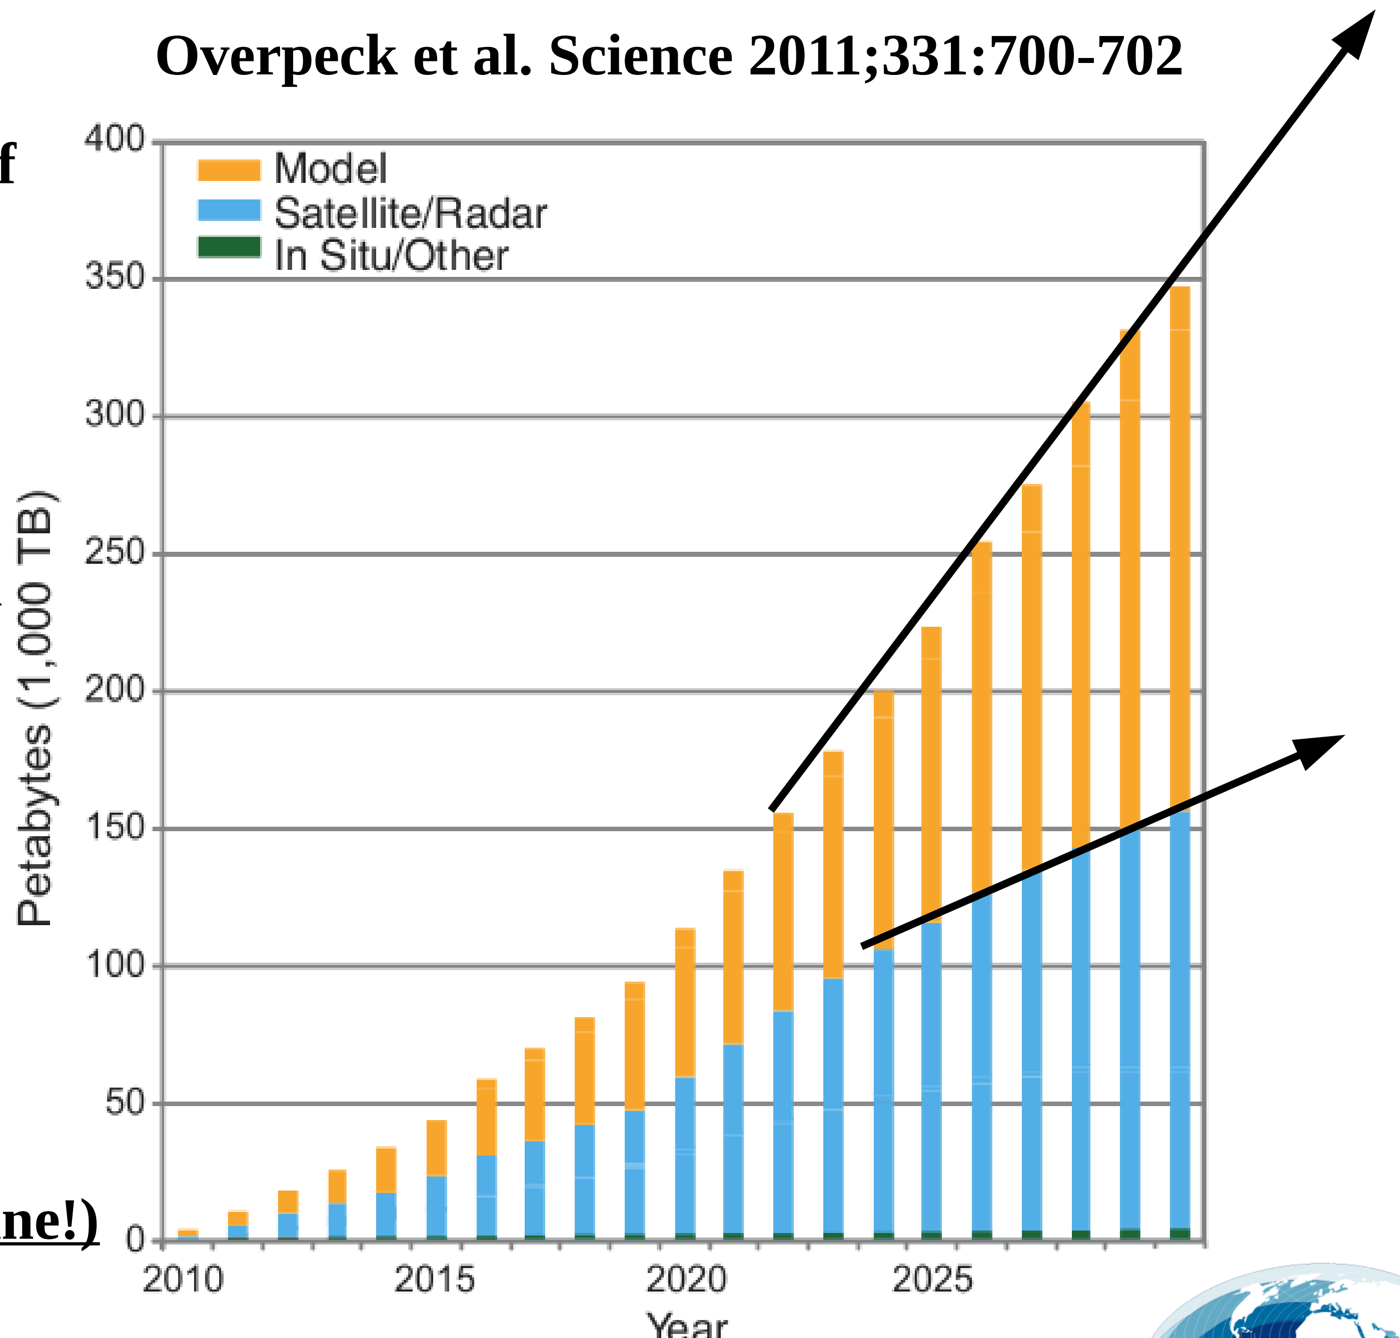

(Their words, not mine!)

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

prodiguer

We might want to write more data

Another way of looking at this:

CMIP5 output at the IPSL was about 3TB / day in average (2 years).

We wanted to have pick writing at 6 TB/day sometime.

Actually can't wrote at full model capacity every day.

- We had to stop simulations from time to time to be able to cope with data production !
- **Archive and network were the real constraints! Already!**

& that's before considering users, usability etc.

Now we know that archive and network are constraints, let's turn to accessibility, affordability, and usability.

A one slide guide to CMIP5 from a data perspective

Fifth Climate Model Intercomparison Project (CMIP5)

World Climate Research Programme WCRP- WGCM
Involves all the major climate modelling centres.

Original Timing:
o(2) PB of requested output from 20+ modelling centres finished early 2010!
Actual Timing?
Years late.

101 experiments
61 model variants
59,000 datasets!
4.5 million files
2 PB in global archive.
Unknown PB locally!

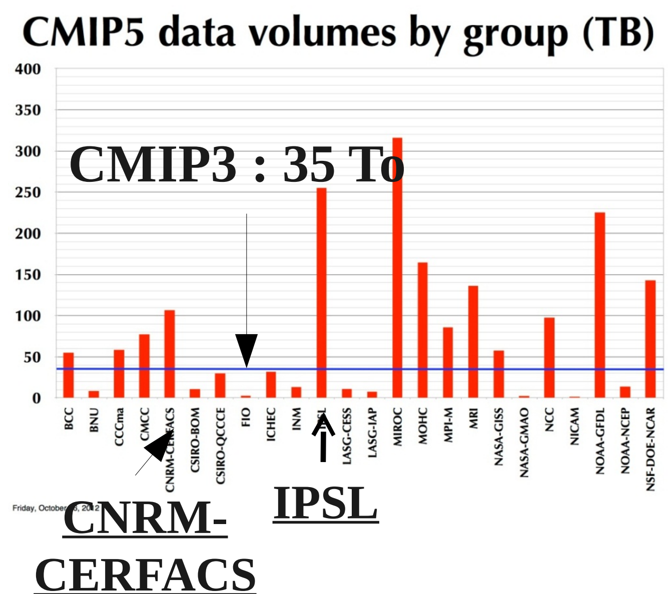

PCMDI-led, Community developed (GO-ESSP) s/w infrastructure for data delivery:
Earth System Grid Federation

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

prodiguer

"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"

A National Strategy for Advancing Climate Modeling, 2012

Why did they think that?

Semantic Analysis: "substantial research effort" "new methods" "computation to data" "rather than trying to download" "frustratingly inaccessible" (to whom?)

...and it's going to get worse,
… more computing ...

« Nor any I should clone » ?

You can buy disk, but, can you populate it?

Network bandwidth limits what we can write (e.g. CMIP5)
Network bandwidth limits what we can deliver for download

- IPSL currently 1 Gb/s
- (Theoretical maximum: 10.5 TB/day. Practical Maximum ~ 0.3 peak)
- It'll go faster (Gilder's Law), fast enough ?

You can buy disk, but, can you afford it?

Decadal Expectations:
Kryder and Kim 2009

Storage Capacity	~ o(x100)
Cost/TB	~ o(/100)
Energy/TB	~ o(/10)
...	
Energy Cost/TB	~ o(x10)

For us (about 1 PB in 2014) at constant funding, we could afford 27 PB (disk) in 2024.

We're not at constant funding. Yet. Are you?

Coming back to National Academy

Analysis !

"substantial research effort"
"new methods"
"computation to data"
"rather than trying to download"
"frustratingly inaccessible" (to whom?)

Download ?

→ What?
(There is a lot of data, how do I choose?)
→ How?
(Where to? How do I make it go fast, multiple files, fault tolerant?)

Computation to Data

Nice idea, but
→ data in multiple places.
→ distributed analysis
→ distributed analysis environments

==>

New methods/research effort

→ What?
- Multiple activities
- ES-DOC

→ How & Computation to Data?
- Avoid duplication?
- EXARCH
- CICLAD

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

prodiguer

… how we interact with data ...

Interacting Data

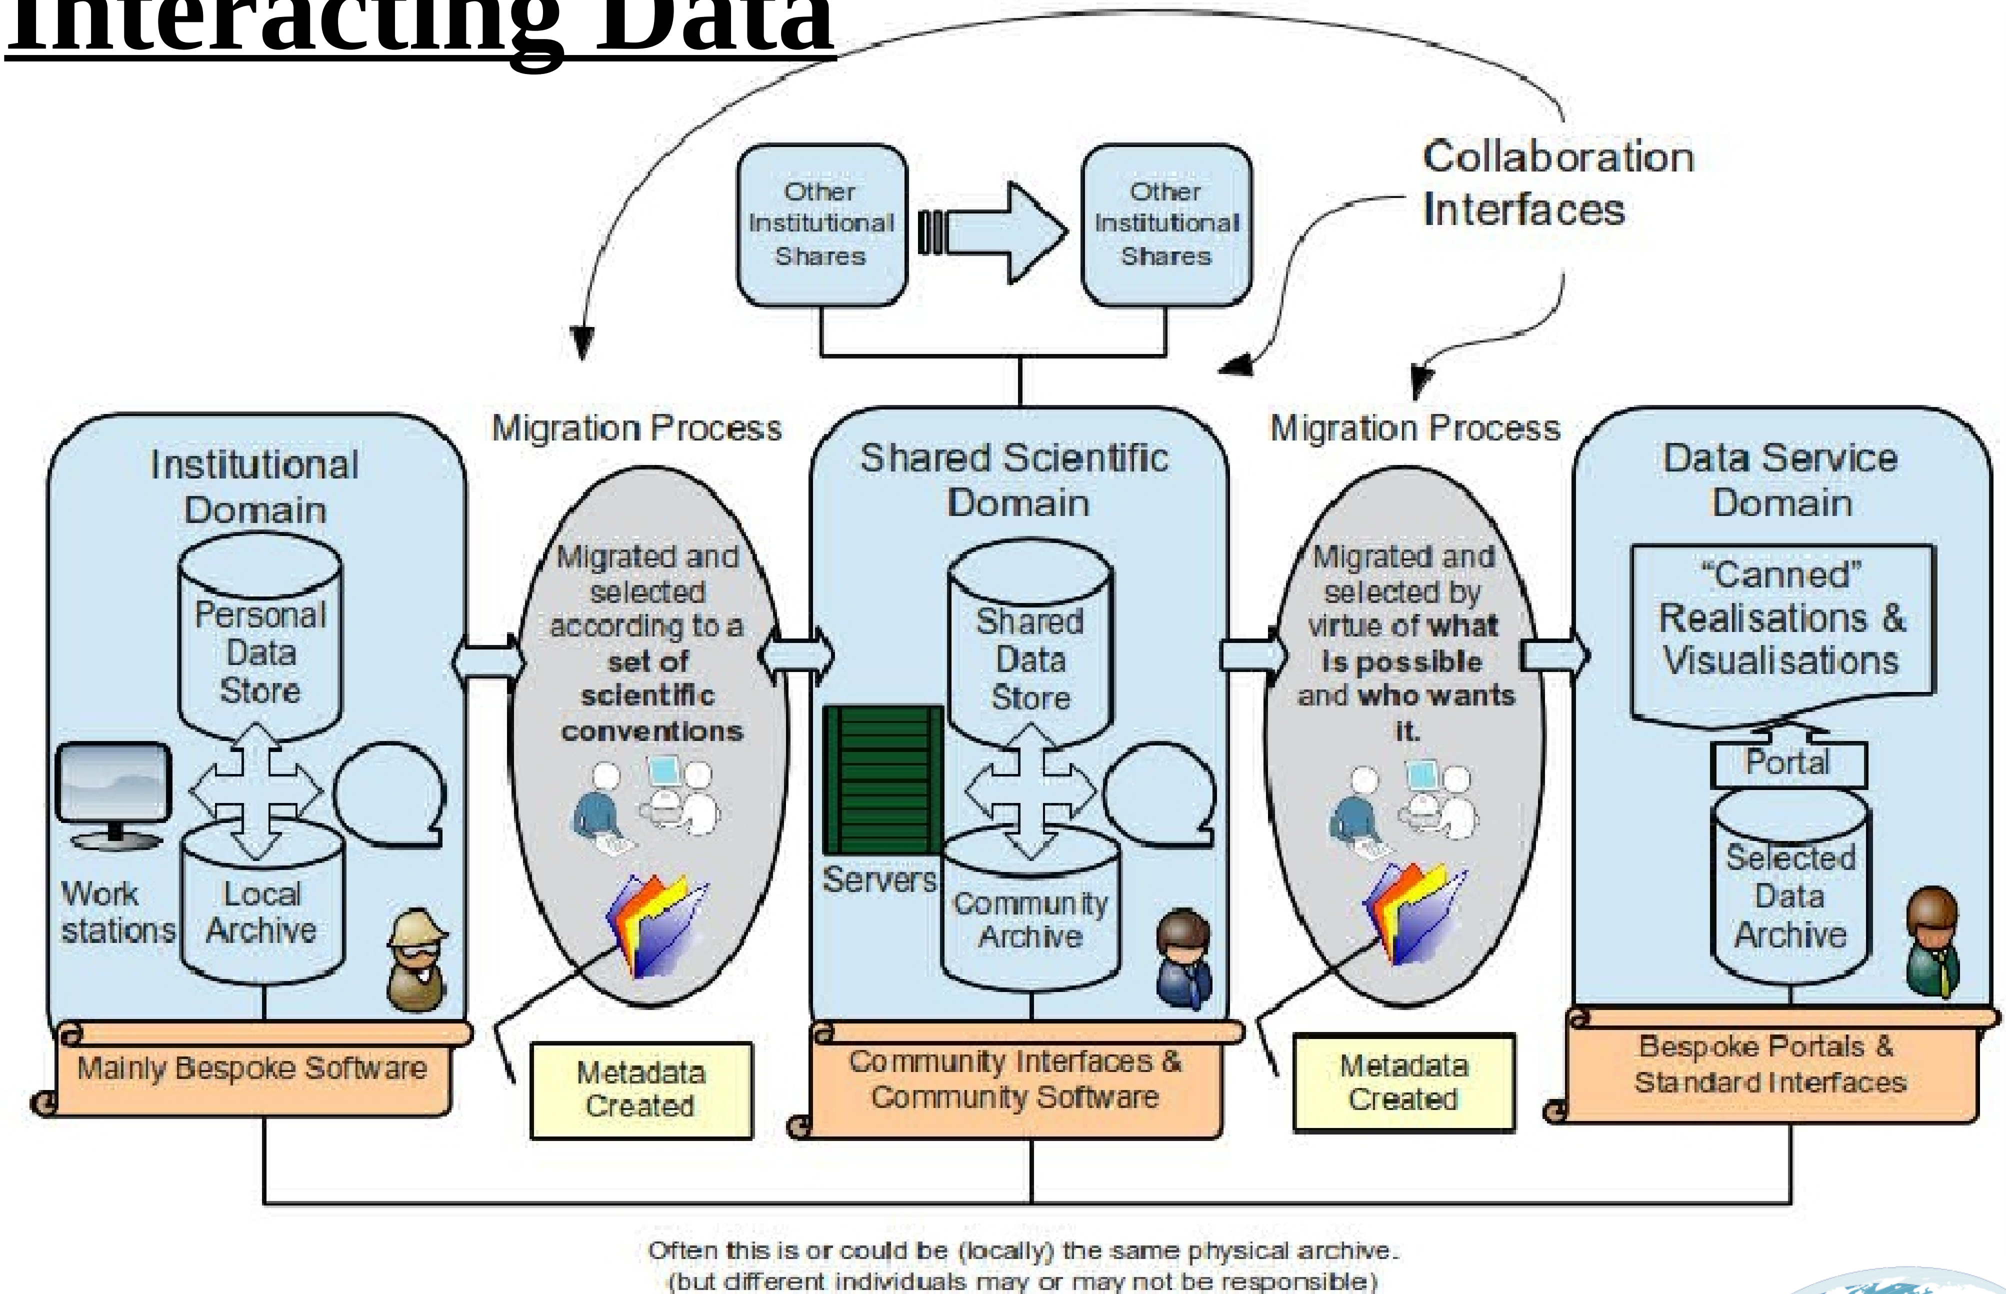

Interacting Data

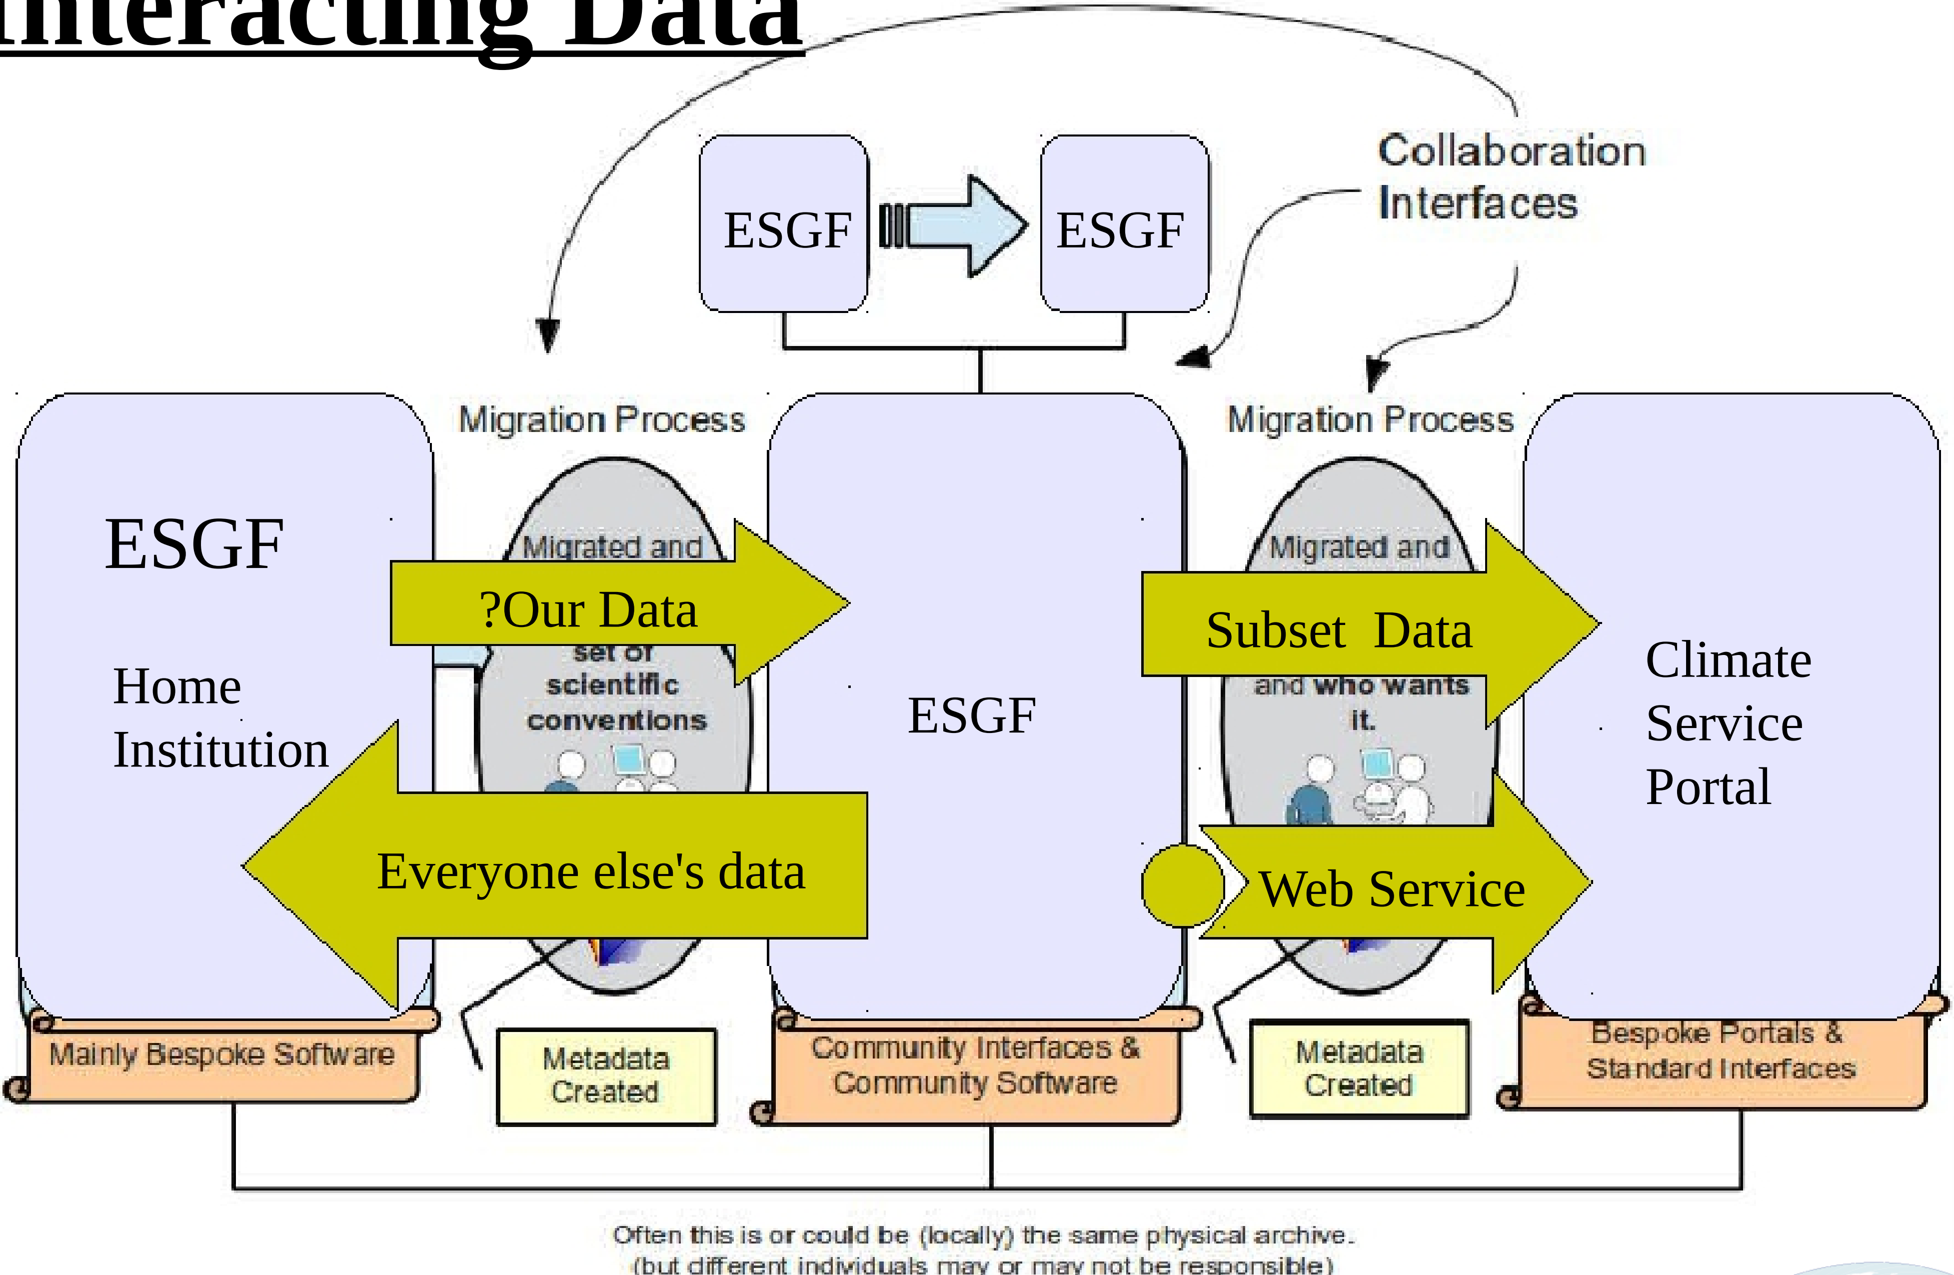

Interacting Data

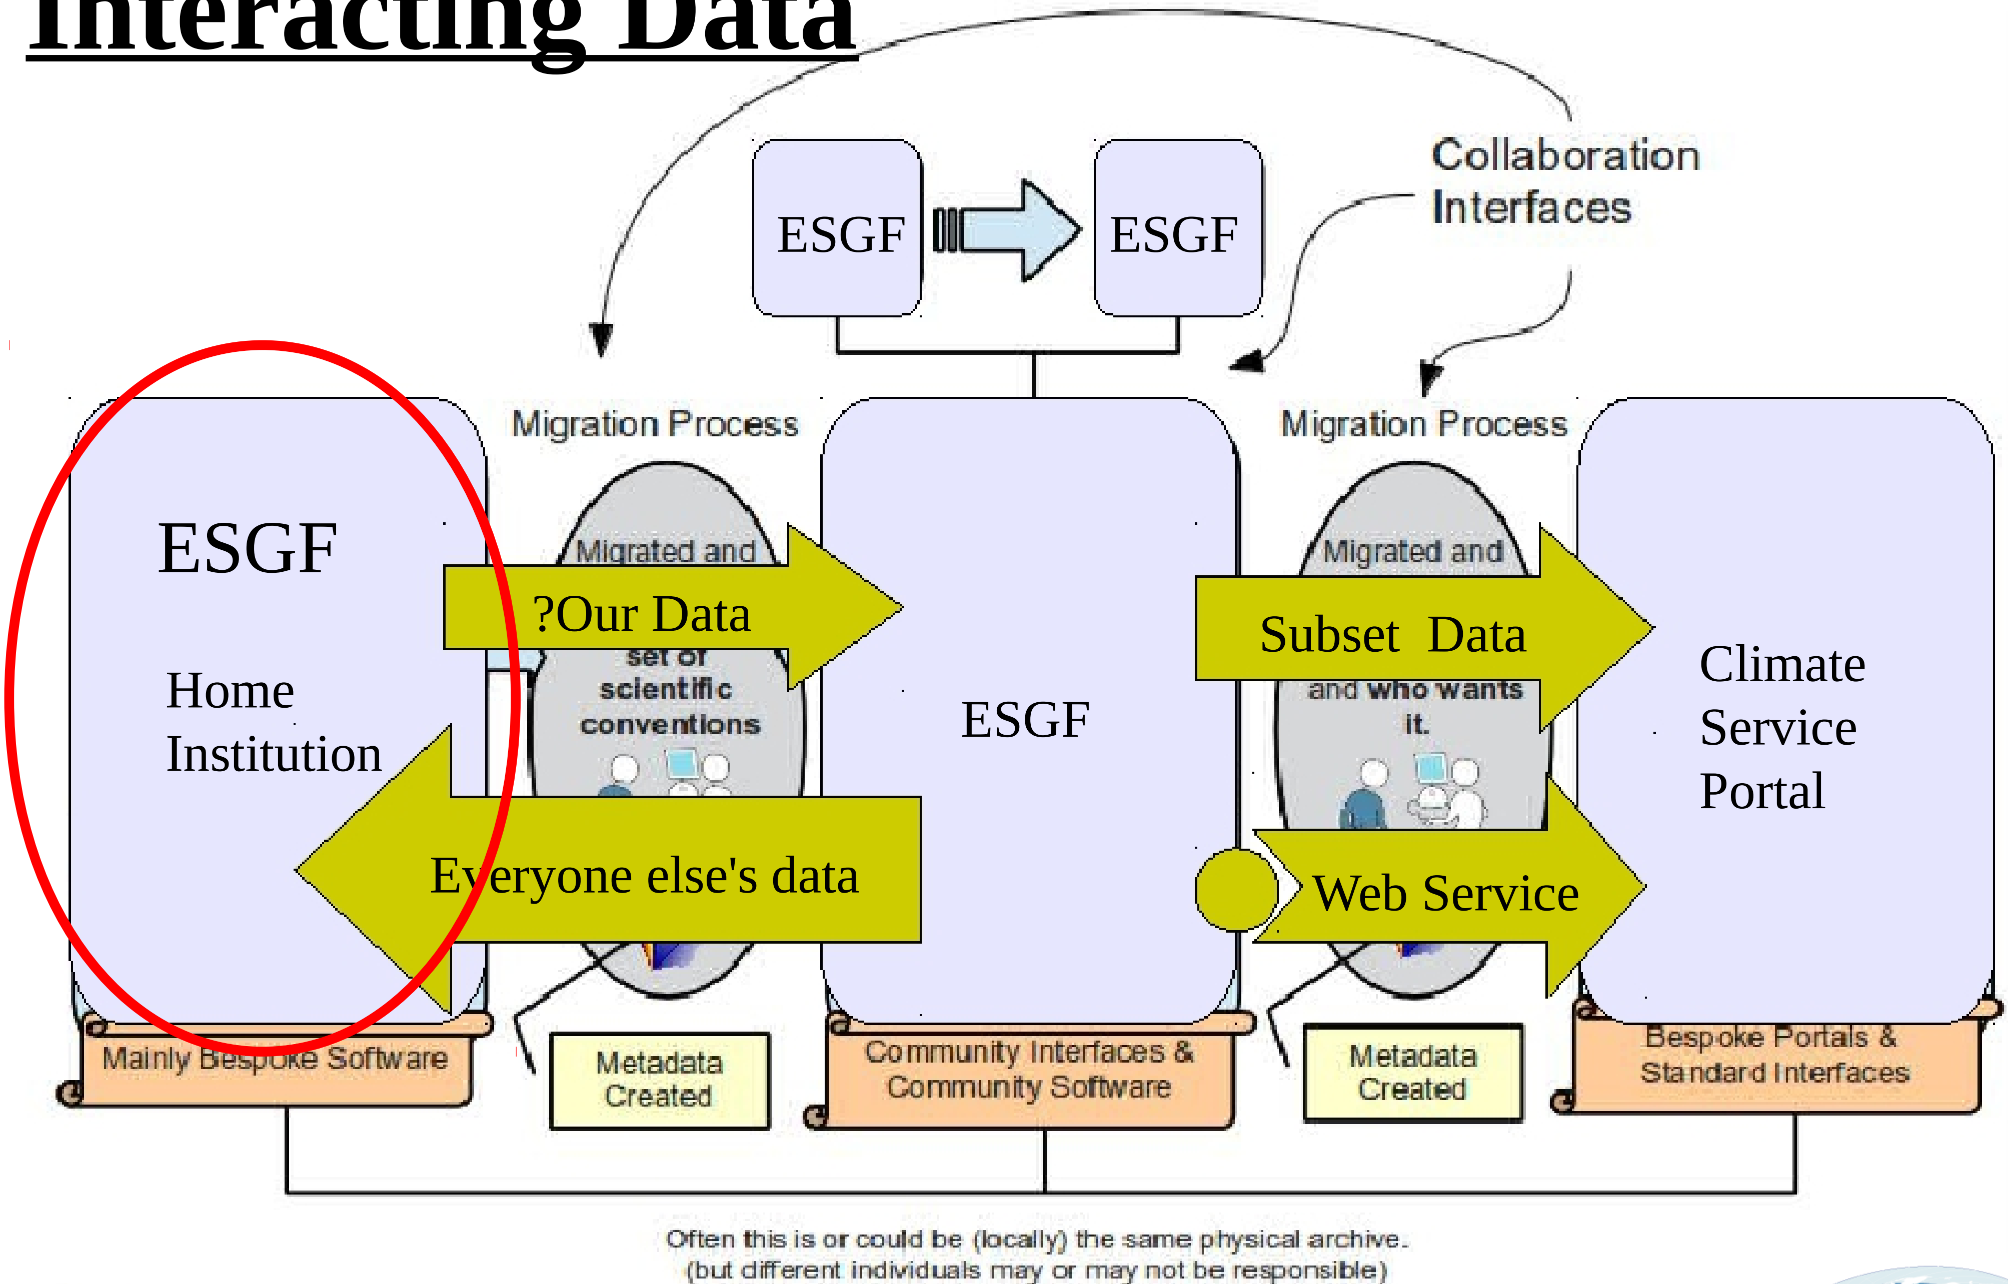

is-enes INFRASTRUCTURE FOR THE EUROPEAN NETWORK FOR EARTH SYSTEM MODELLING

prodiguer

My/Your Data Environment

At your home institution, you:

→ Have (some) control over your software environment

- Favourite packages, e.g. IDL
- Familiar Linux

→ Can buy/arrange more storage / compute on varying time-scales … can optimise ...

→ Are responsible for deleting / preserving your own data

→ Are likely to be duplicating data others have already downloaded *in your own institution* … let alone within a larger collaboration.

We all like control!

We all like the (illusion?) that we can scale our resources as necessary.

We all lose/destroy/duplicate data.

Most of us do our HPC remotely.

Some of us do our analysis remotely. Why not more of us?

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

prodiguer

Interacting Data

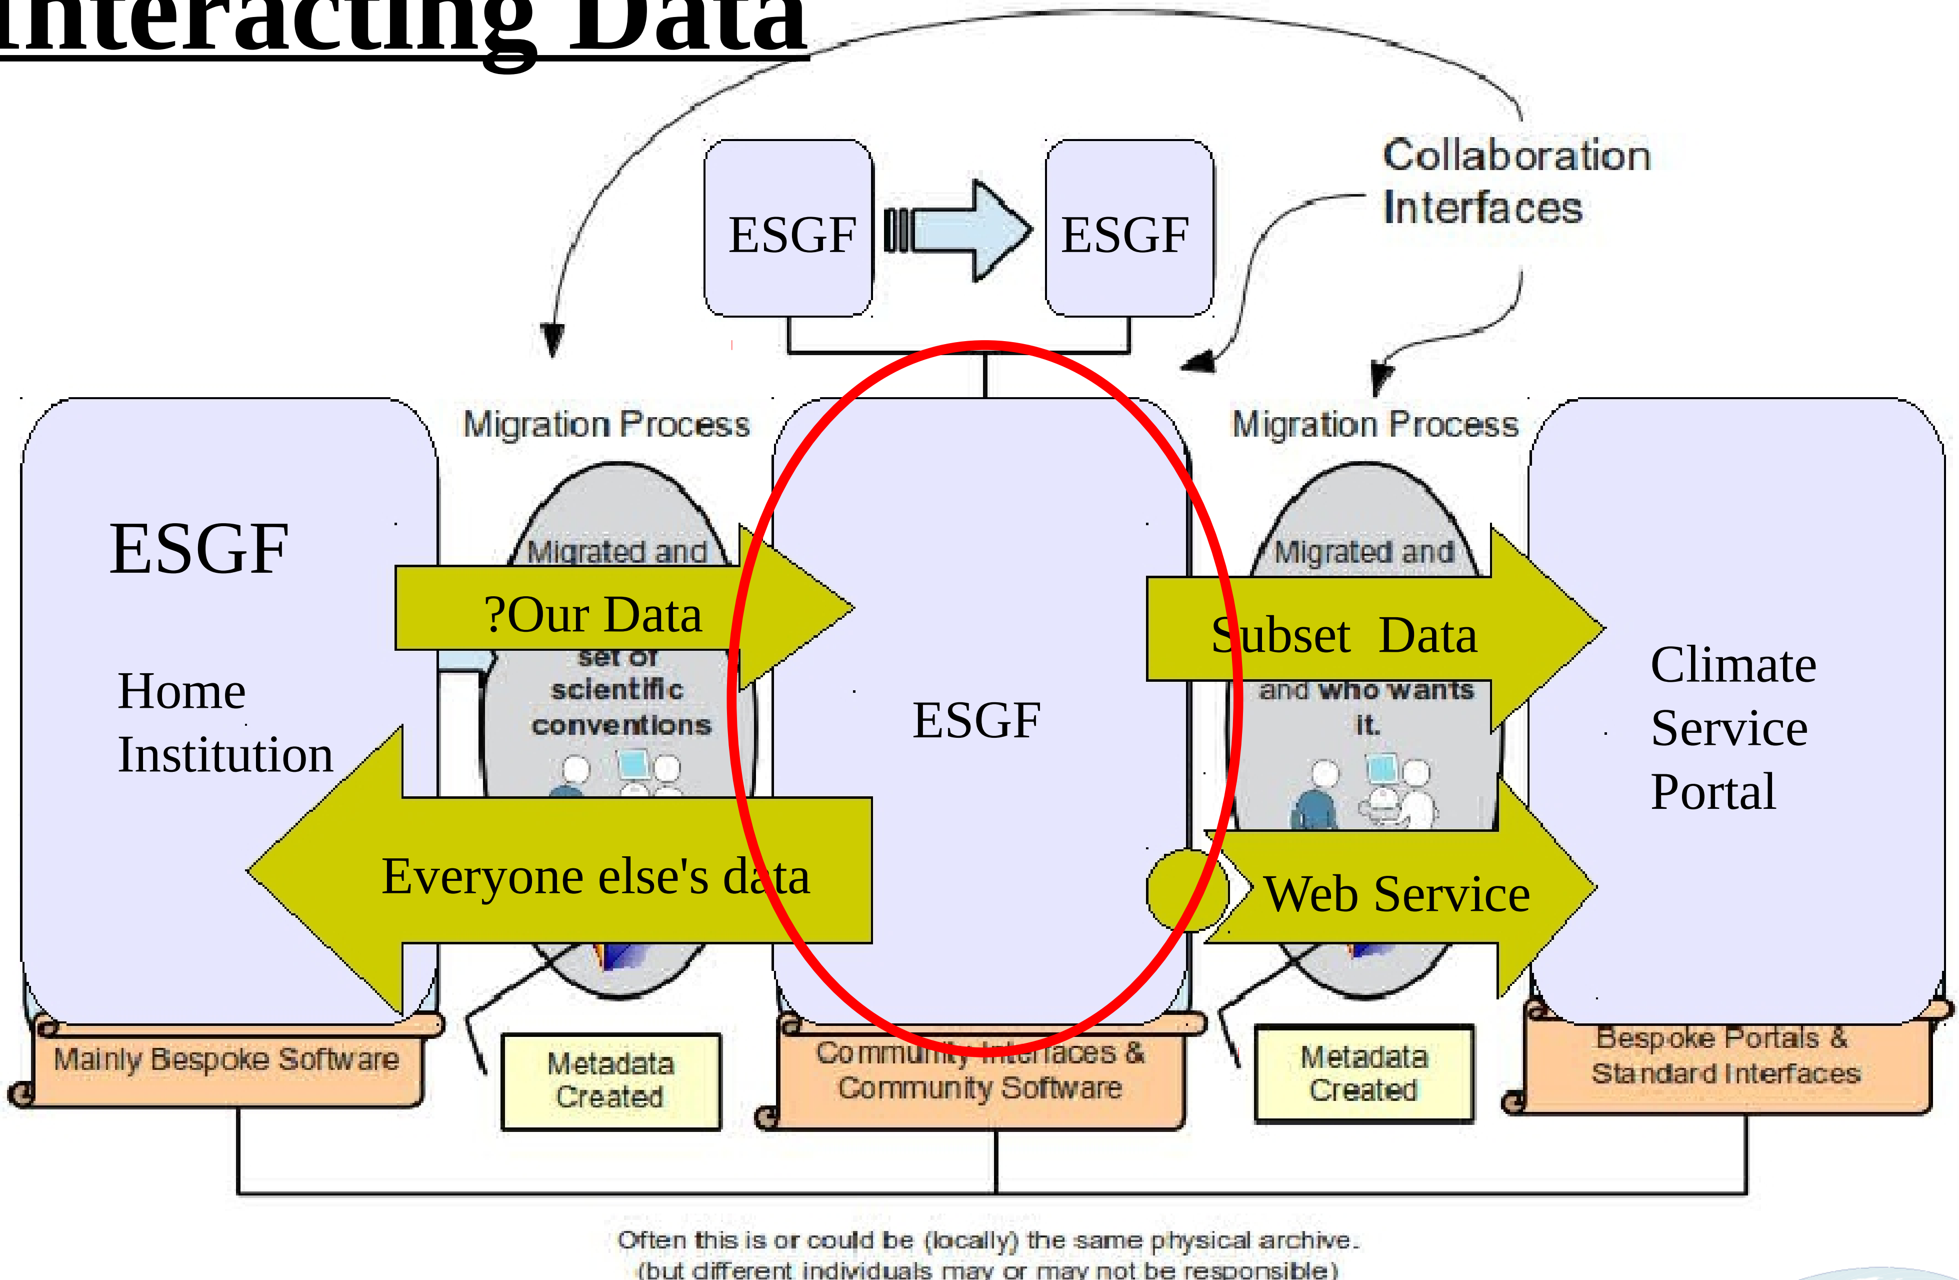

A one slide guide to CMIP5 from a data perspective

Fifth Climate Model Intercomparison Project (CMIP5)

World Climate Research Programme
WCRP- WGCM
Involves all the major climate modelling centres.

Original Timing:
o(2) PB of requested output from 20+ modelling centres finished early 2010!
Actual Timing?
Years late.

101 experiments
61 model variants
59,000 datasets!
4.5 million files
2 PB in global archive.
Unknown PB locally!

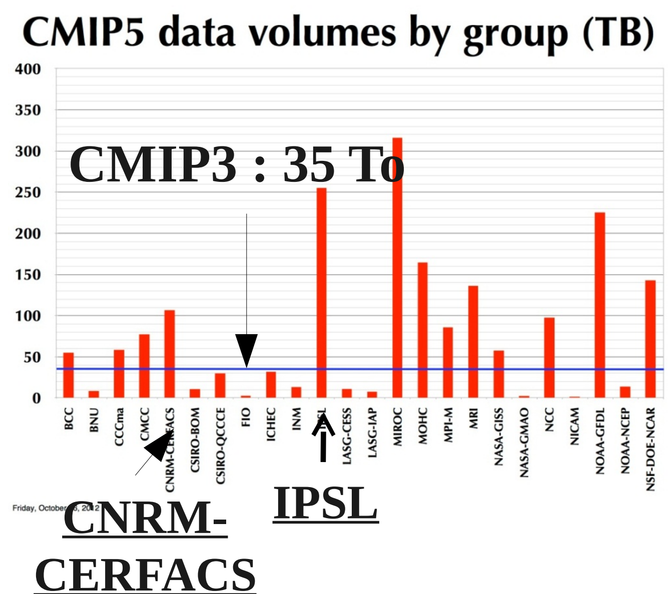

PCMDI-led,
Community developed (GO-ESSP)
s/w infrastructure for data delivery:
Earth System Grid Federation

Worldwide distributed system

The **Earth System Grid Federation (ESGF)** is a multi-agency, international collaboration of persons and institutions working together to build an open source software infrastructure for the management and analysis of Earth Science data on a global scale

- Software development and project management: ANL, ANU, **BADC**, **CMCC**, **DKRZ**, ESRL, GFDL, GSFC, JPL, **IPSL**, ORNL, LLNL (lead), PMEL, …
- Operations: tens of data centers across Asia, Australia, Europe and North America

IPCC AR5 distribution

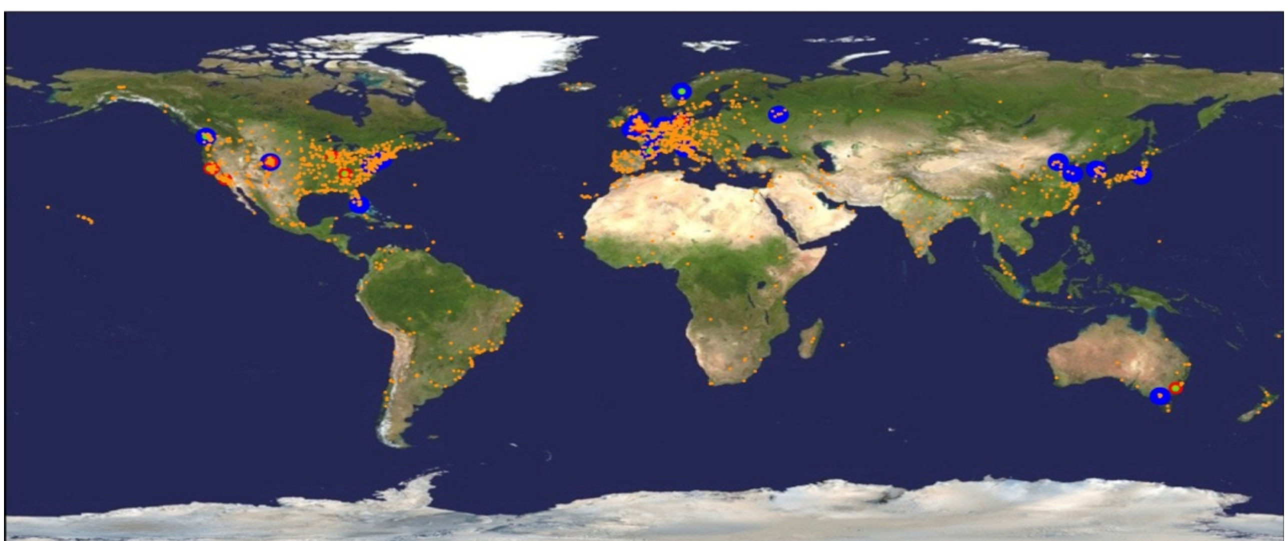

Modeling centers (24) Gateways (9) Nodes (14) Users (1000s)

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

prodiguer

Worldwide distributed system

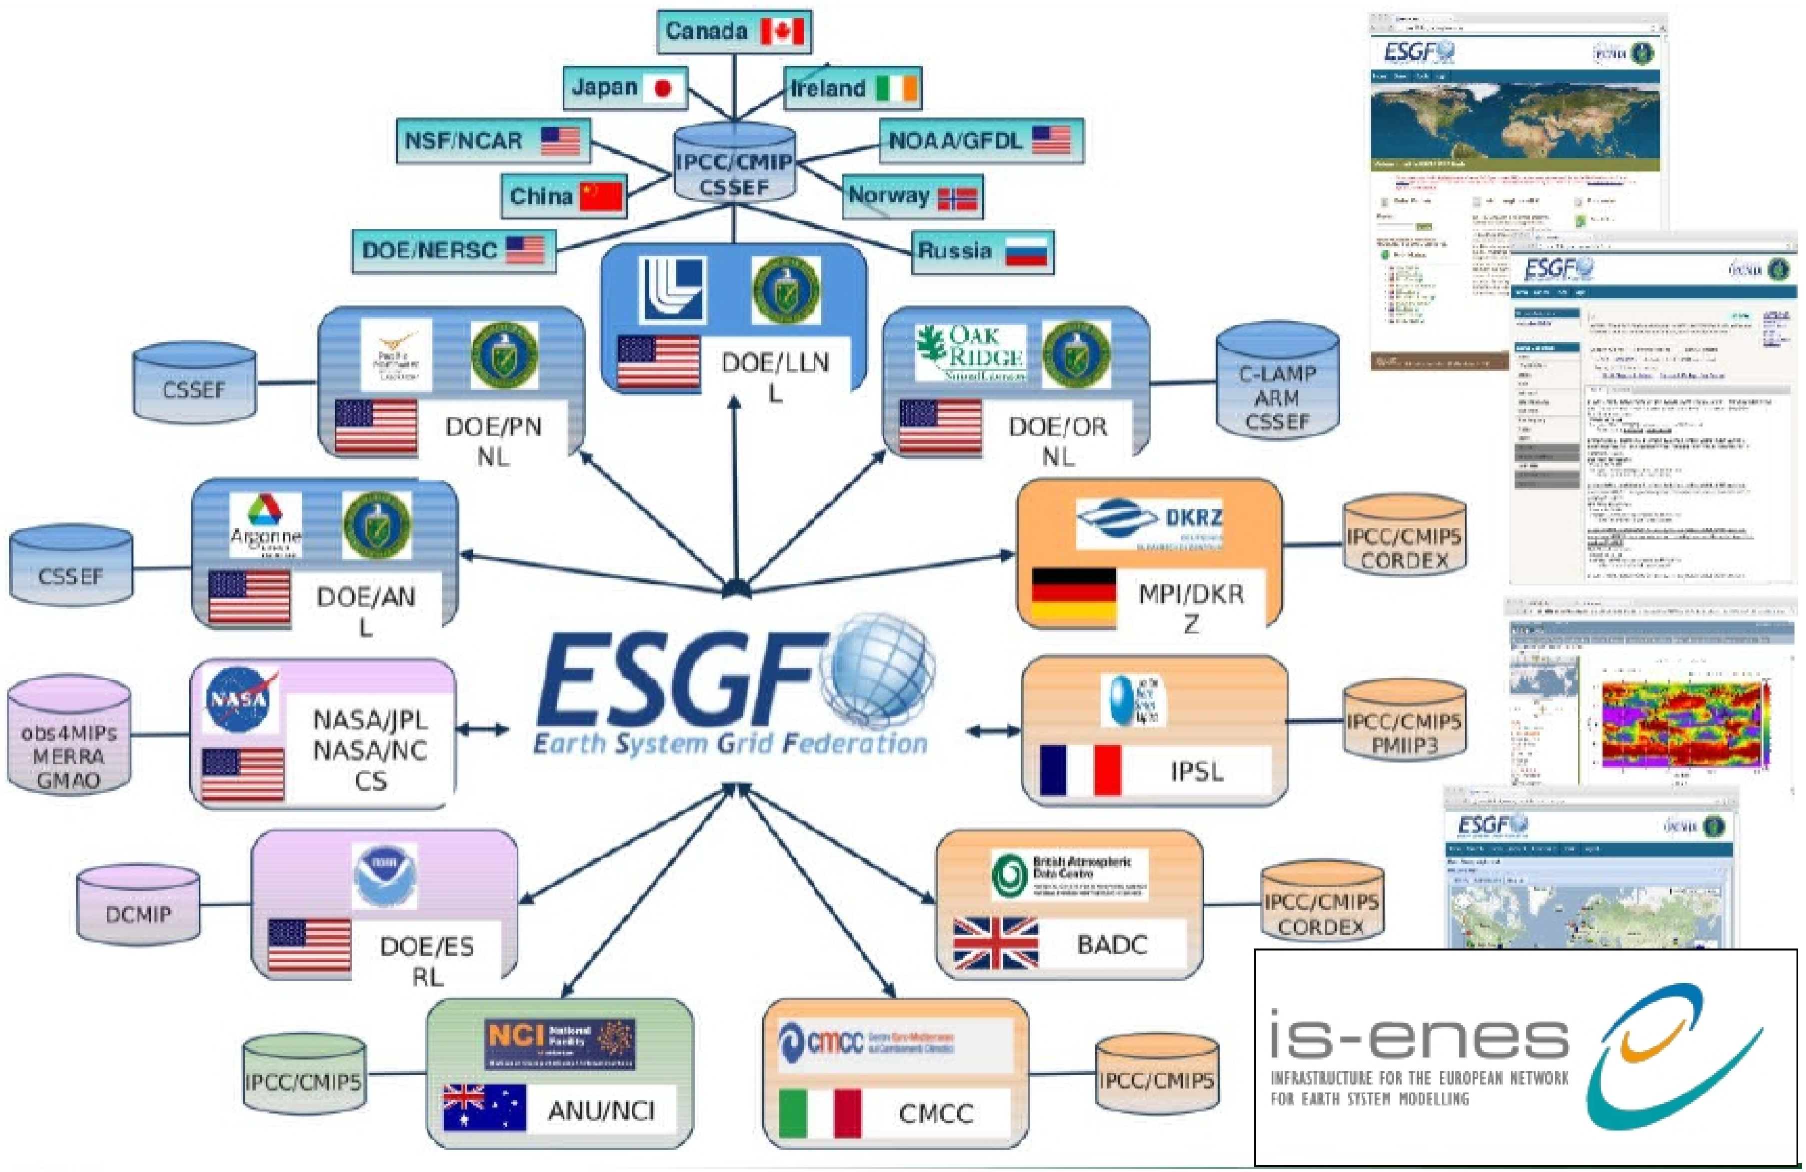

Institut Pierre Simon Laplace — METEO FRANCE Toujours un temps d'avance — CERFACS — IDRIS — TGCC — CINES Centre Informatique National de l'Enseignement Supérieur

A one slide guide to CMIP5 from a data perspective

Fifth Climate Model Intercomparison Project (CMIP5)

World Climate Research Programme WCRP- WGCM
Involves all the major climate modelling centres.

Original Timing:
o(2) PB of requested output from 20+ modelling centres finished early 2010!
Actual Timing?
Years late.

101 experiments
61 model variants
59,000 datasets!
4.5 million files
2 PB in global archive.
Unknown PB locally!

CMIP5 data volumes by group (TB)

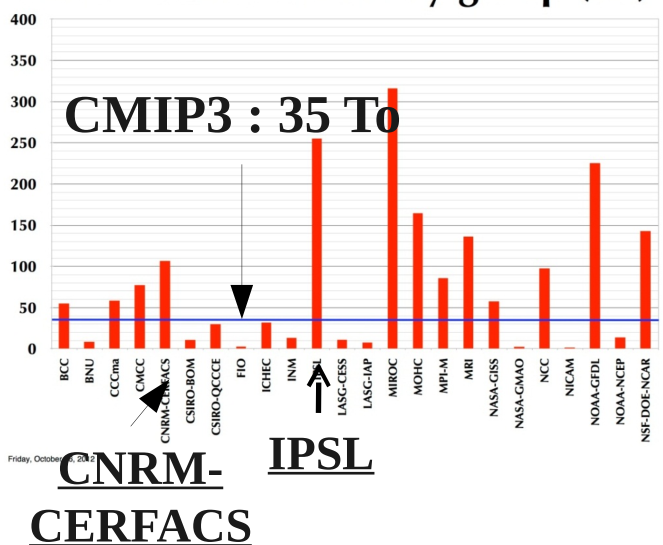

PCMDI-led, Community developed (GO-ESSP) s/w infrastructure for data delivery:
Earth System Grid Federation

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

prodiguer

A climate simulation

http://earthsystemcog.org/projects/es-doc-models/

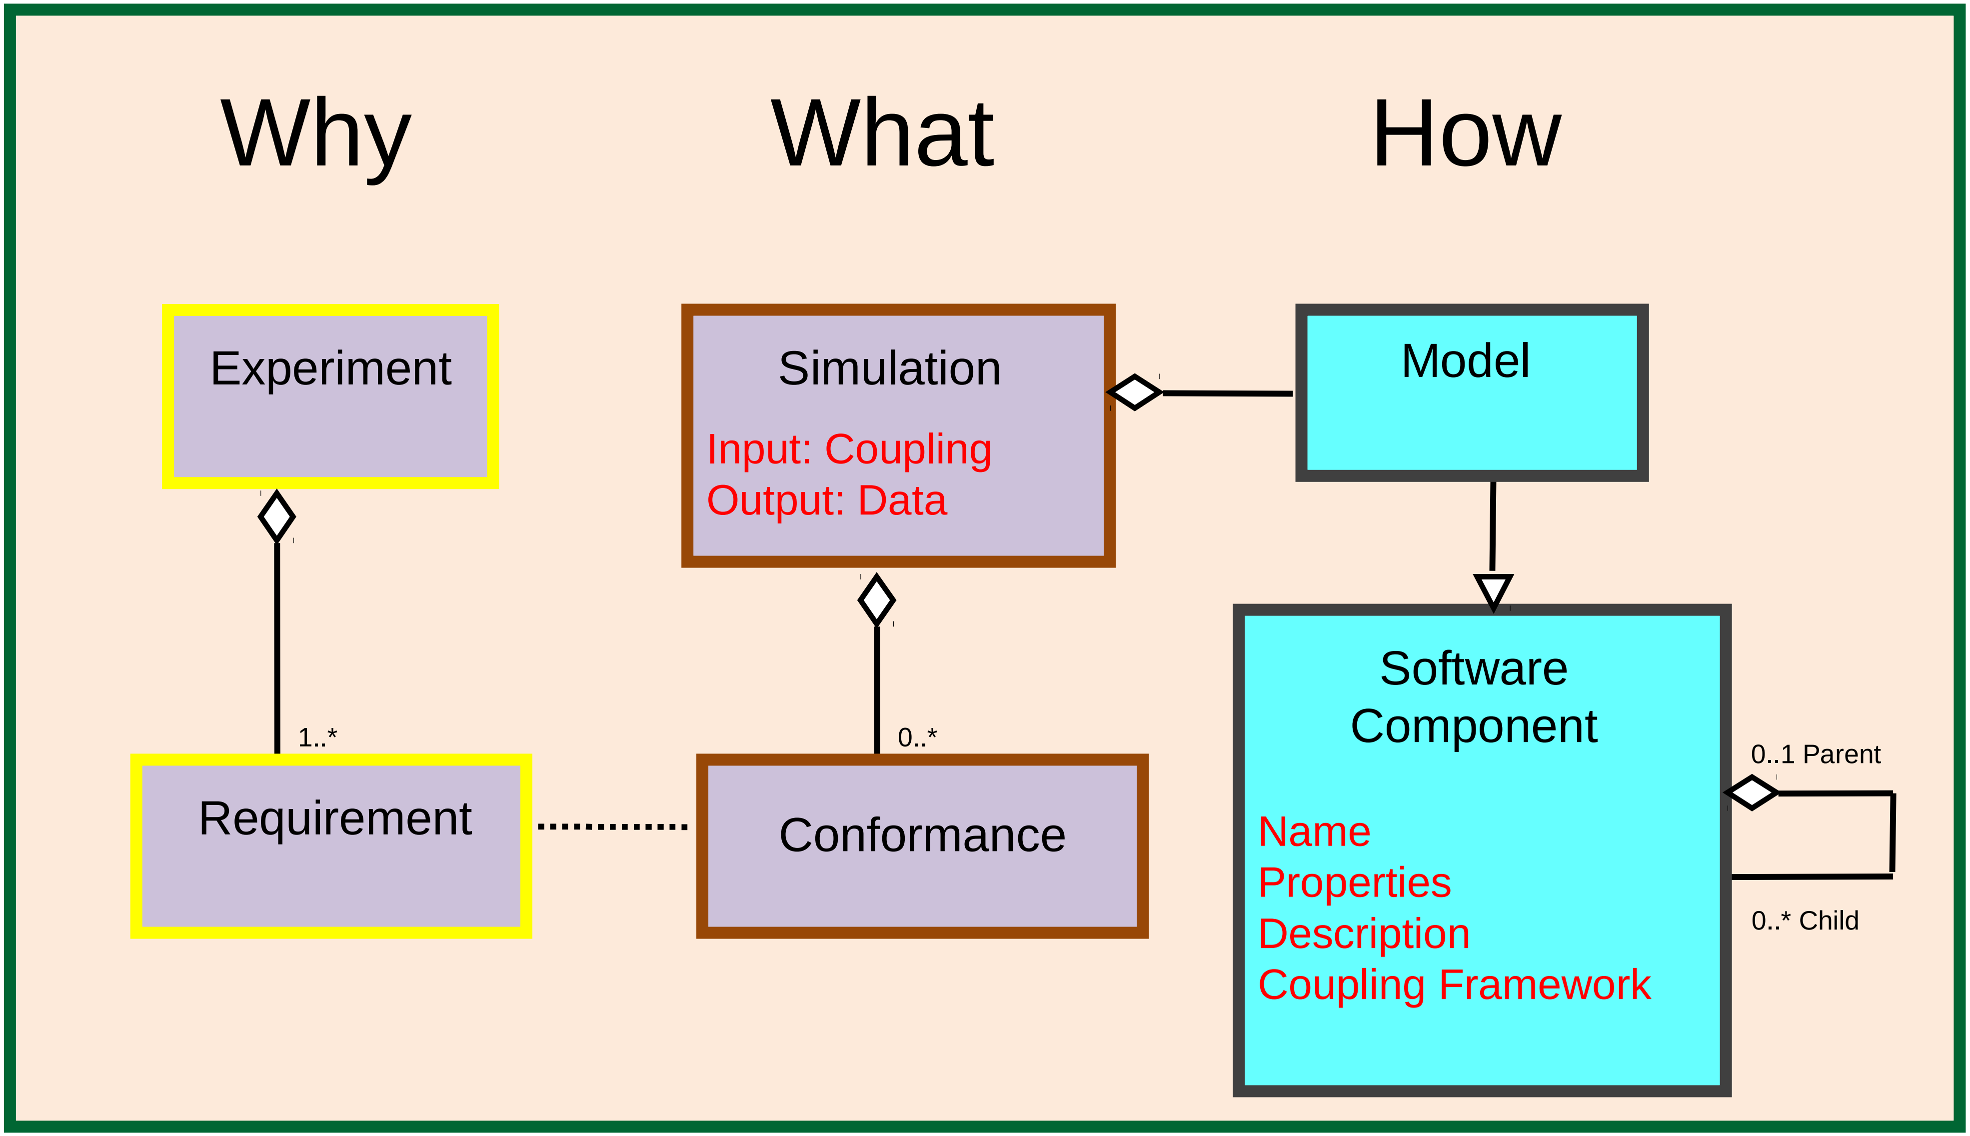

CMIP5 – Browse Metadata

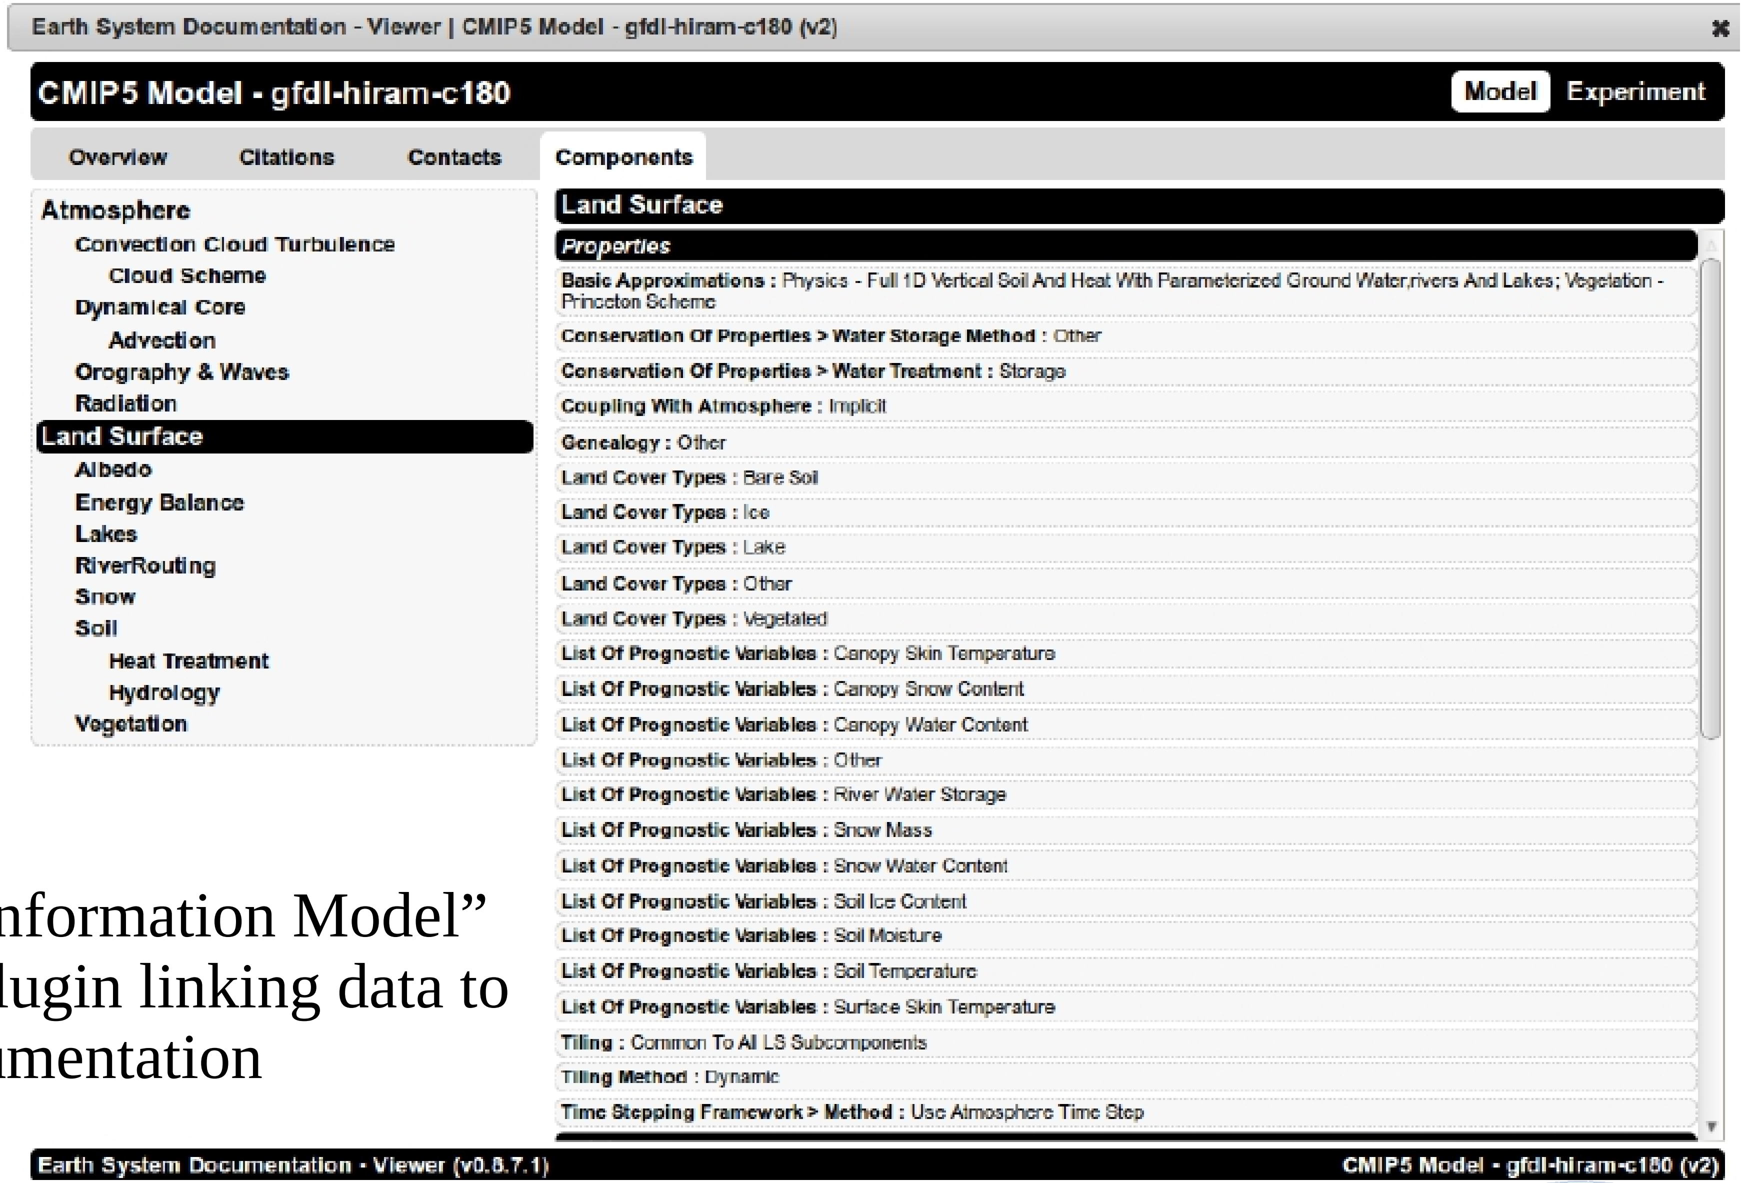

ES-DOC
“Common Information Model”
Javascript plugin linking data to remote documentation

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

prodiguer

CMIP5 - Comparing and contrasting

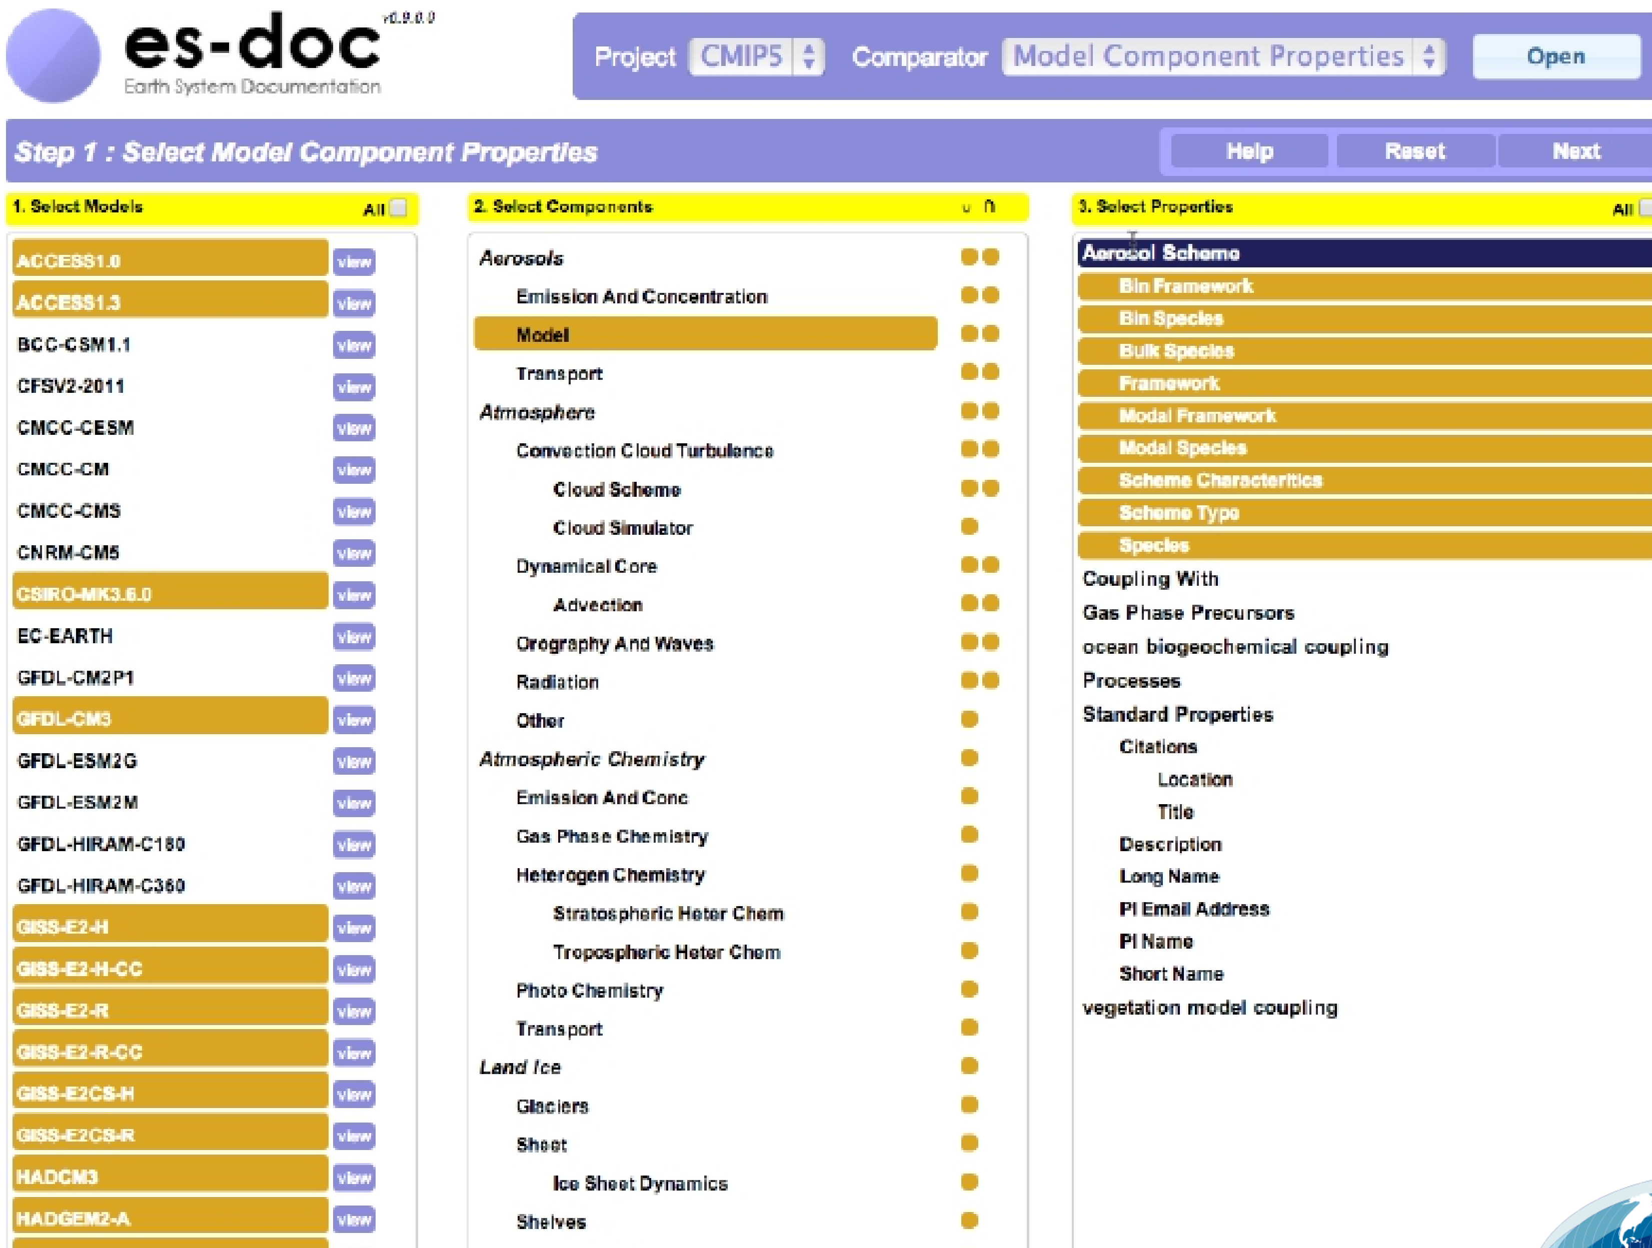

is-enes
INFRASTRUCTURE FOR THE EUROPEAN NETWORK FOR EARTH SYSTEM MODELLING

prodiguer

CMIP archive (aka disk) Metadata

Two levels: the way we bundle the files (the layout on disk), and the content of the files (the layout of metadata) within them!

The Layout:
Directory Reference Syntax (DRS)

Key to supporting tools which manage the (millions of files), and provide services. Key to a RESTful API ...

The Files
Constrained CF (Climate Model Output Rewriter, CMOR, and tables).

- Define layout *in* files: so tools can manipulate the file contents.

- key role for netCDF climate forecast conventions **and** additional (project) metadata resquirements. That get just the bundle « you »

… it's all about supporting automation! Key to usability!

So we've just seen how ESGF works, how much of that will apply at exascale?

New methods/research effort

→ Computation to Data?
Avoid duplication?
EXARCH
CICLAD

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

prodiguer

ExArch

G8 funded research project aiming at 10 years horizon:
ExArch: Climate analytics on distributed exascale data archives.
STFC (NCAS), UCLA, U. Toronto, DKRZ, Princeton U., IPSL, CMCC

Web processing services
Query syntax
Common information model
Processing operators and quality control
Scientific diagnostics
EO data for model evaluation
Grid computing

"Bring the computing to the data"
→ need faith in the data (QC)
→ need to know what the data is (metadata)
→ exploit existing operators (tooling)

Requires coordination between data repositories ...

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

prodiguer

High Performance Data (HPD) Analysis Environment

Mutualized

Jointly *delivered* by
→ IPSL laboratories.
Joint *users* (initially):
→ IPSL community
Joint *users* (target):
→ French Academic community

Analysis capabilities

Environmental Data
Compute Service
Web Service Provision for :
→ Climate Science
→ Earth Observation
→ Environmental studies

Access services to ESGF System

Users don't have to find, download, and keep up to date the data they need

CMIP5, CORDEX
Reanalysis,
Obs4MIPs
PMIP
...

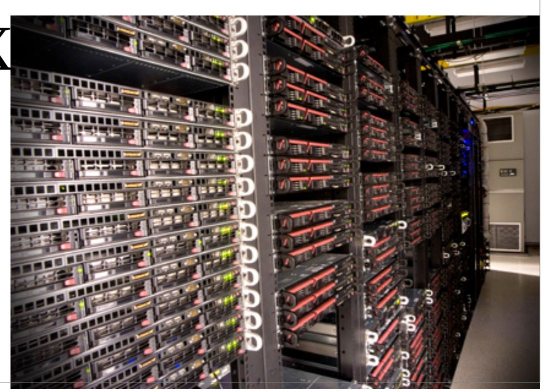

Big DATA Platform

Collaboration Environment
→ Access to Curated Archive.
→ Large shared "Group Work Spaces"
→ climate analysis enabled system
→ + 1 PB of high performance disk coupled to hundreds of cores configured for analysis

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

prodiguer

Summary

→ We might find the path to exascale is constrained by storage as much as by compute!

→ Collaboration depends on metadata, with many different levels involved.

→ The dependency on metadata will only grow as the scale of our archives increase.

→ We will need not only to move computation to data, but aggregate our data collections

  → which requires dedicated "HPD" platforms like CICLAD

→ It is likely that those who cannot conform to shared standards will (should) be left working on their own, unable to enter their data into federated and shared systems.

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

prodiguer