Why Cloud? Earth Systems Science Perspective

or

Data Driven Science

Bringing Computation to the Data

Whether that data started life in an instrument or a computer!

Bryan N Lawrence

Outline

- The Big Picture: Communities and Infrastructure
- Background Trends: Output Data Growth
- Our Cloud Approach
- Summary

Direct Numerical Simulation

Primarily mathematical representation of a complex system of processes

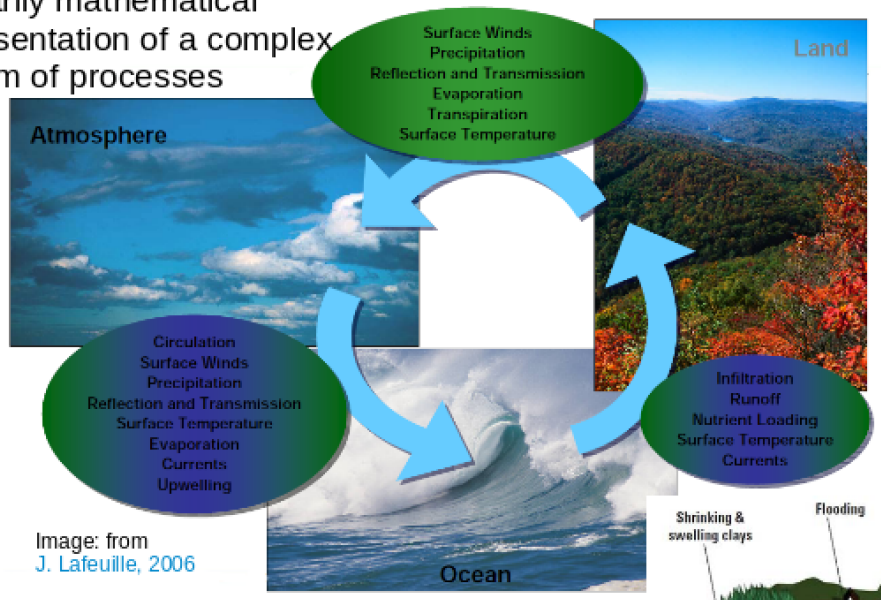

Image: from J. Lafeuille, 2006

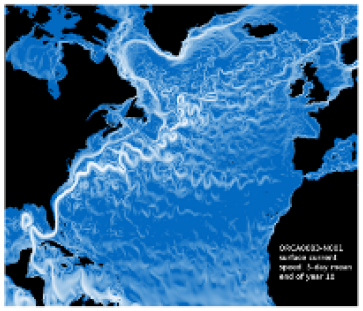

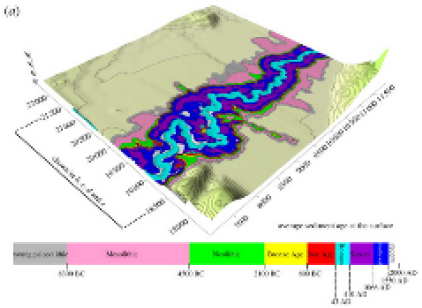

Coulthard and Van De Wiel IDoi: 10.1098/rsta.2011.0597

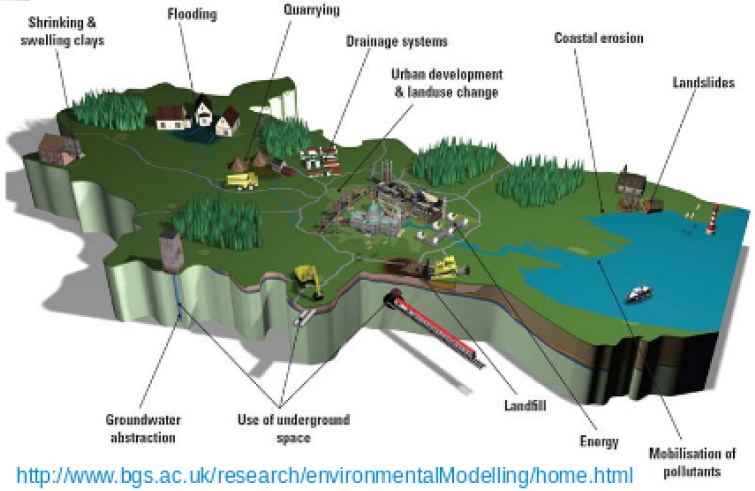

http://www.bgs.ac.uk/research/environmentalModelling/home.html

We want to observe and simulate the world at ever higher resolution! More complexity!

Increasing complexity

FAR:1990
SAR:1995
TAR:2001
AR4:2007
AR5:2013

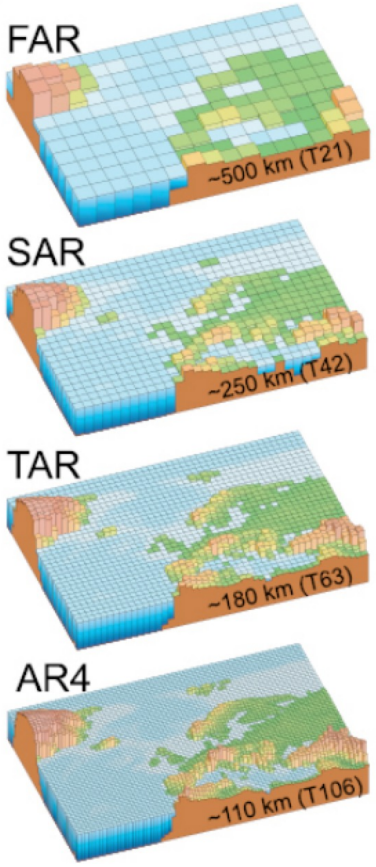

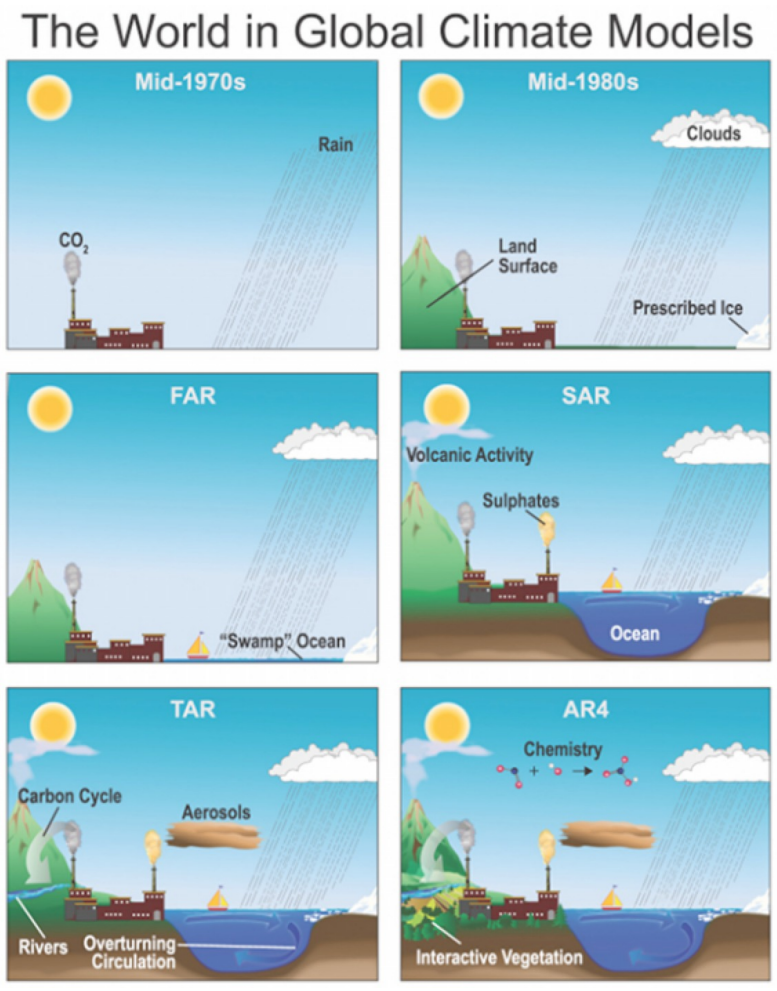

ESGF

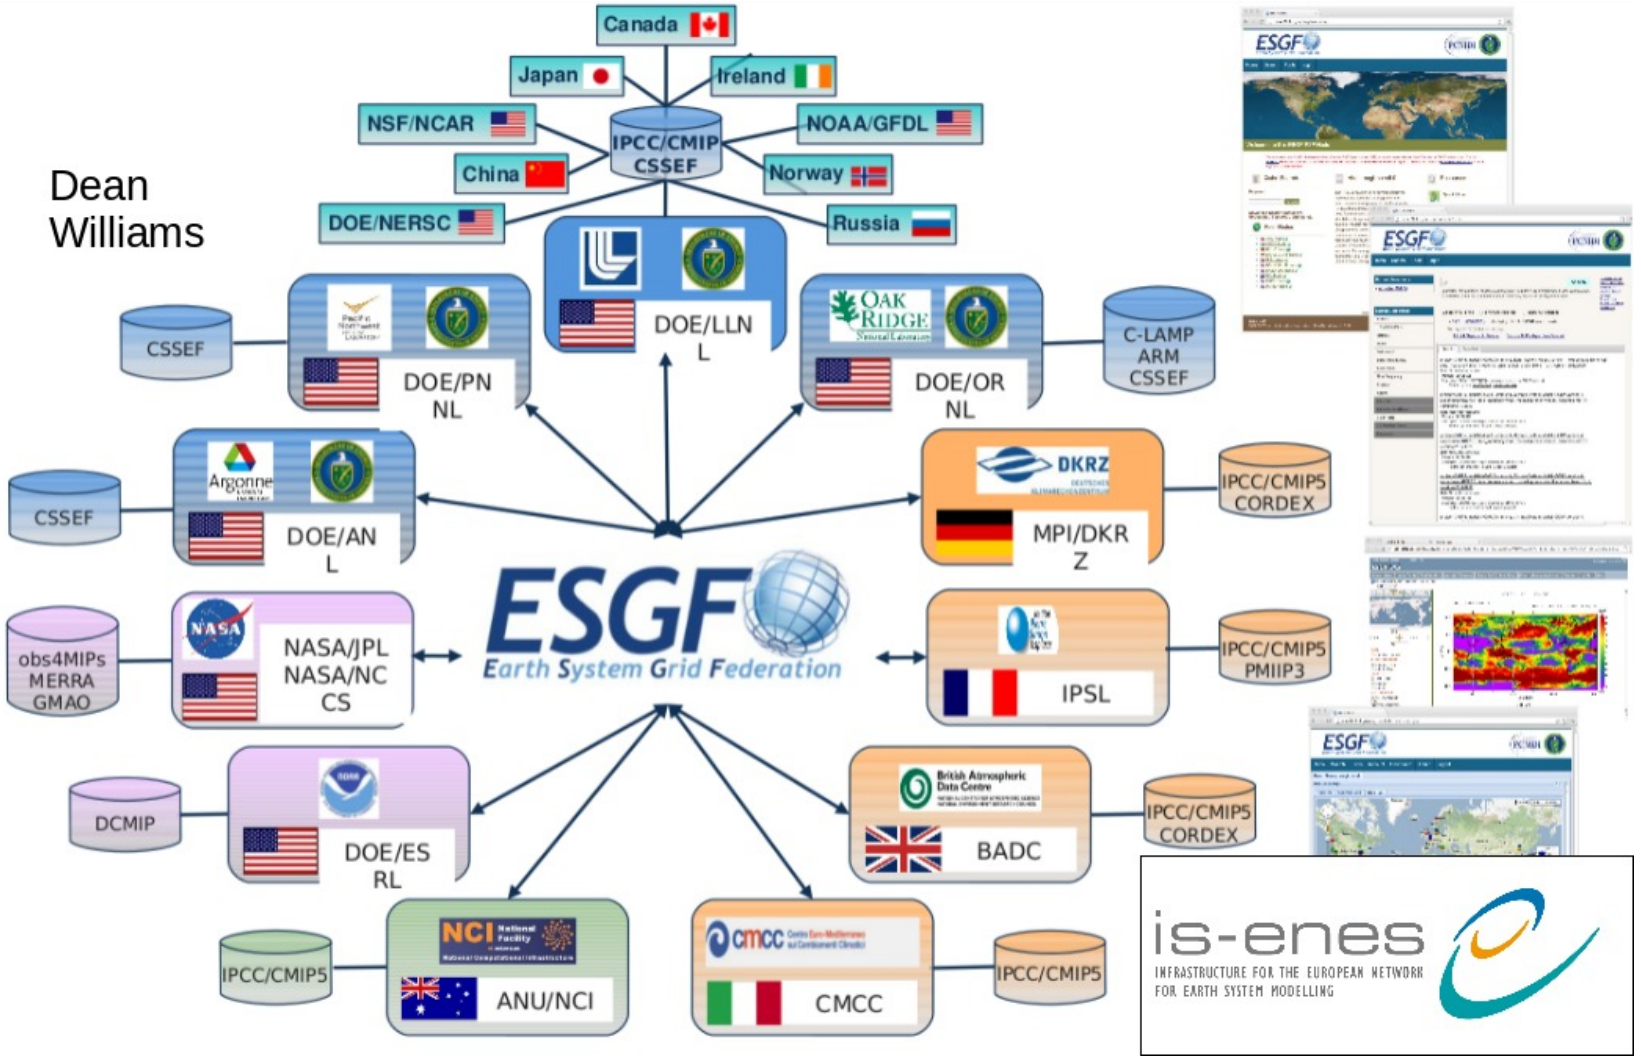
Dean Williams
Canada
Japan
Ireland
NSF/NCAR
IPCC/CMIP CSSEF
NOAA/GFDL
China
Norway
DOE/NERSC
Russia
DOE/LLNL
CSSEF
DOE/PNNL
DOE/ORNL
C-LAMP ARM CSSEF
CSSEF
DOE/ANL
DKRZ
MPI/DKRZ
IPCC/CMIP5 CORDEX
ESGF
Earth System Grid Federation
obs4MIPs MERRA GMAO
NASA/JPL NASA/NCCS
IPSL
IPCC/CMIP5 PMIIP3
DCMIP
DOE/ESRL
BADC
IPCC/CMIP5 CORDEX
IPCC/CMIP5
NCI National Facility
ANU/NCI
cmcc
CMCC
IPCC/CMIP5
is-enes
INFRASTRUCTURE FOR THE EUROPEAN NETWORK FOR EARTH SYSTEM MODELLING

Communities

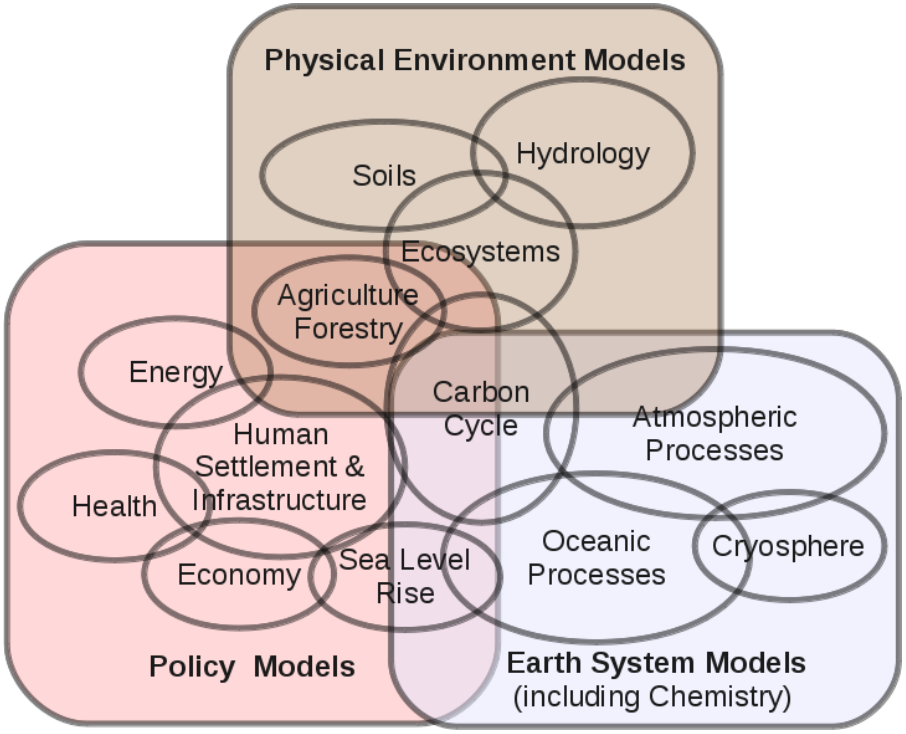

Many interacting communities, each with their own software, compute environments etc.

Figure adapted from Moss et al, 2010

Where is this going?

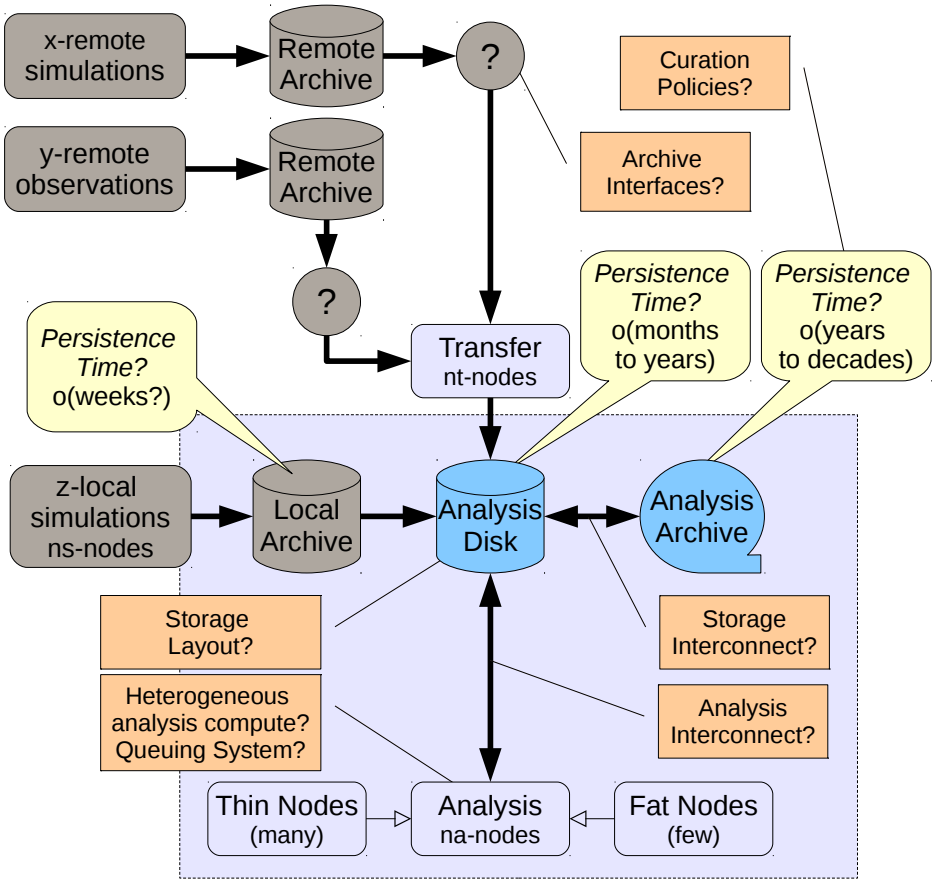

- (Potentially) many different remote simulation sources. How long can the data remain at source?
- Interesting problems moving the data to a common location?
- How long can the data reside on disk at the analysis location? What about in the archive?
- How should we best organise the data?
- What are the best ways to organise analysis compute?
- What are the best ways to address analysis interconnect and I/O bandwidth?

Sharing

Science across scales

Lots of interacting communities

Lots of infrastructure

New sorts of infrastructure

Can we share infrastructure?

Global Data Archival

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're not a climate scientist.

(BNL: Even if you are?)

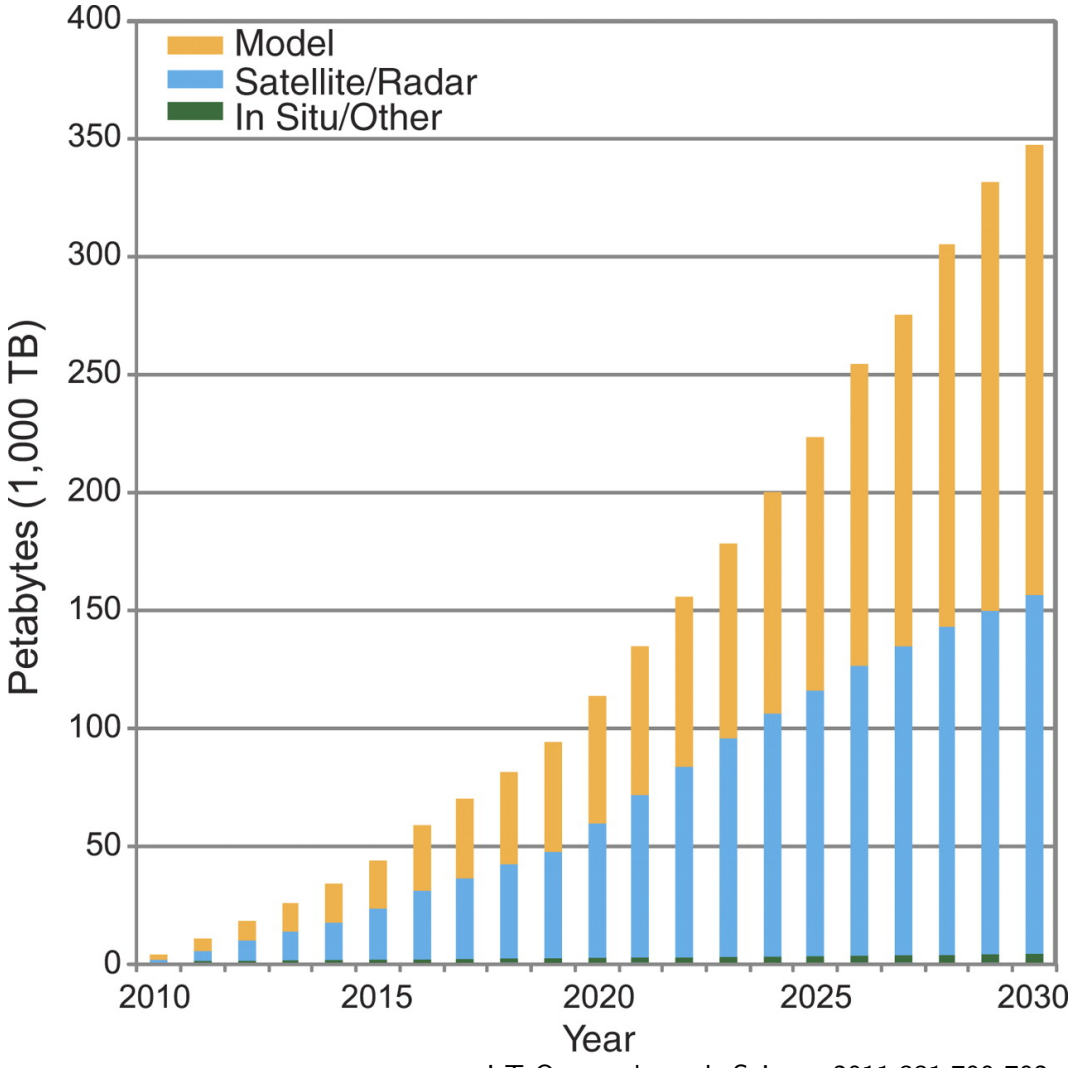

J T Overpeck et al. Science 2011;331:700-702

Institutional - NCAR

Storage, and power for storage, will dominate NCAR's compute budget within a few years! (Rich Loft, 2014).

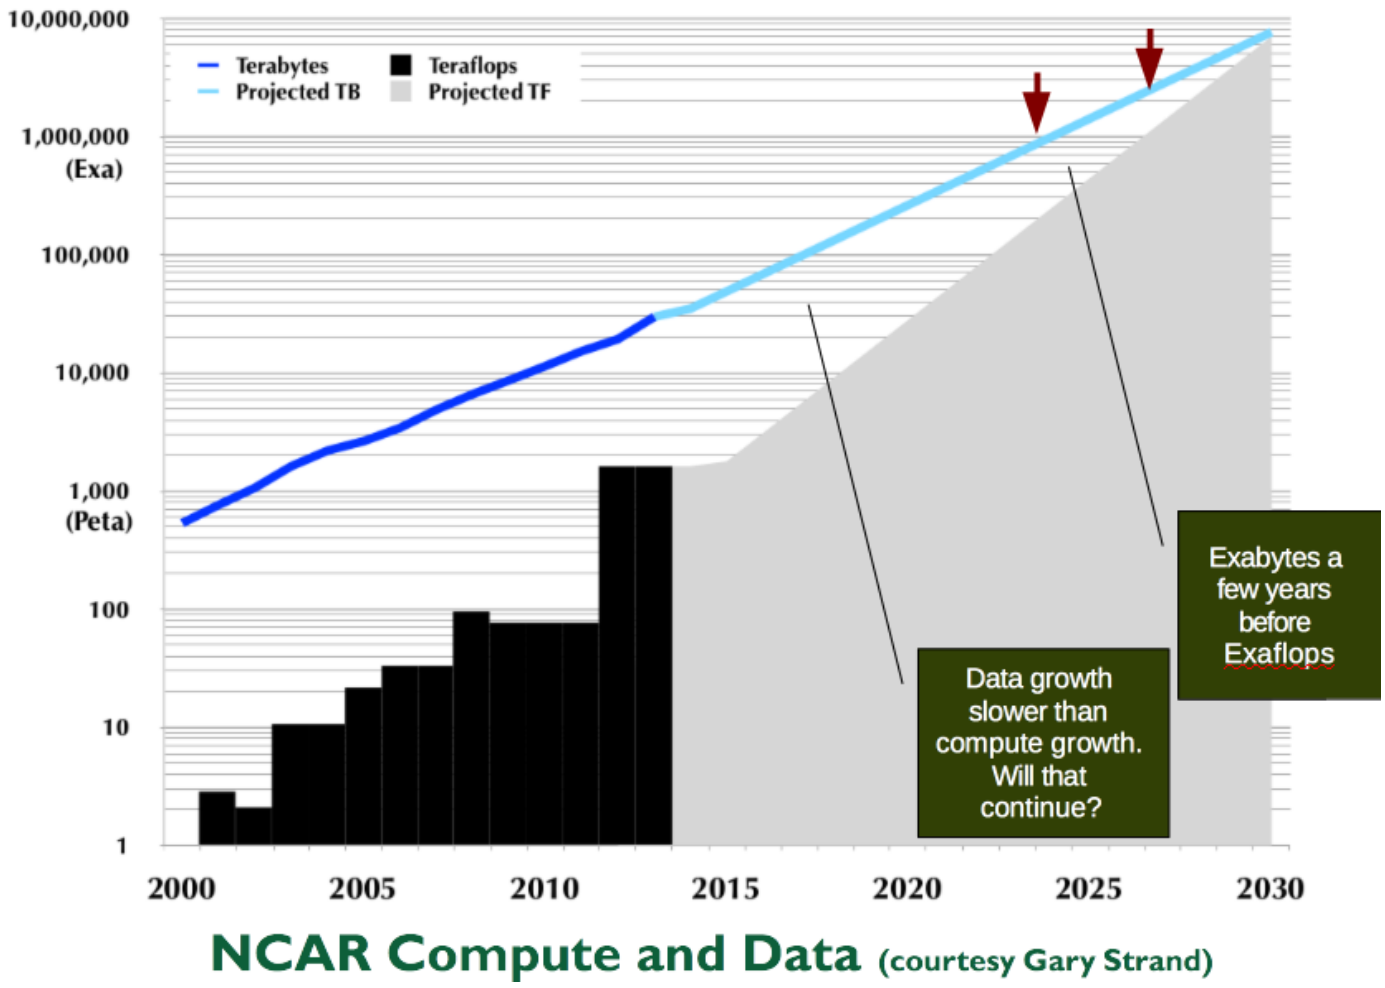

Institutional - STFC and CEDA

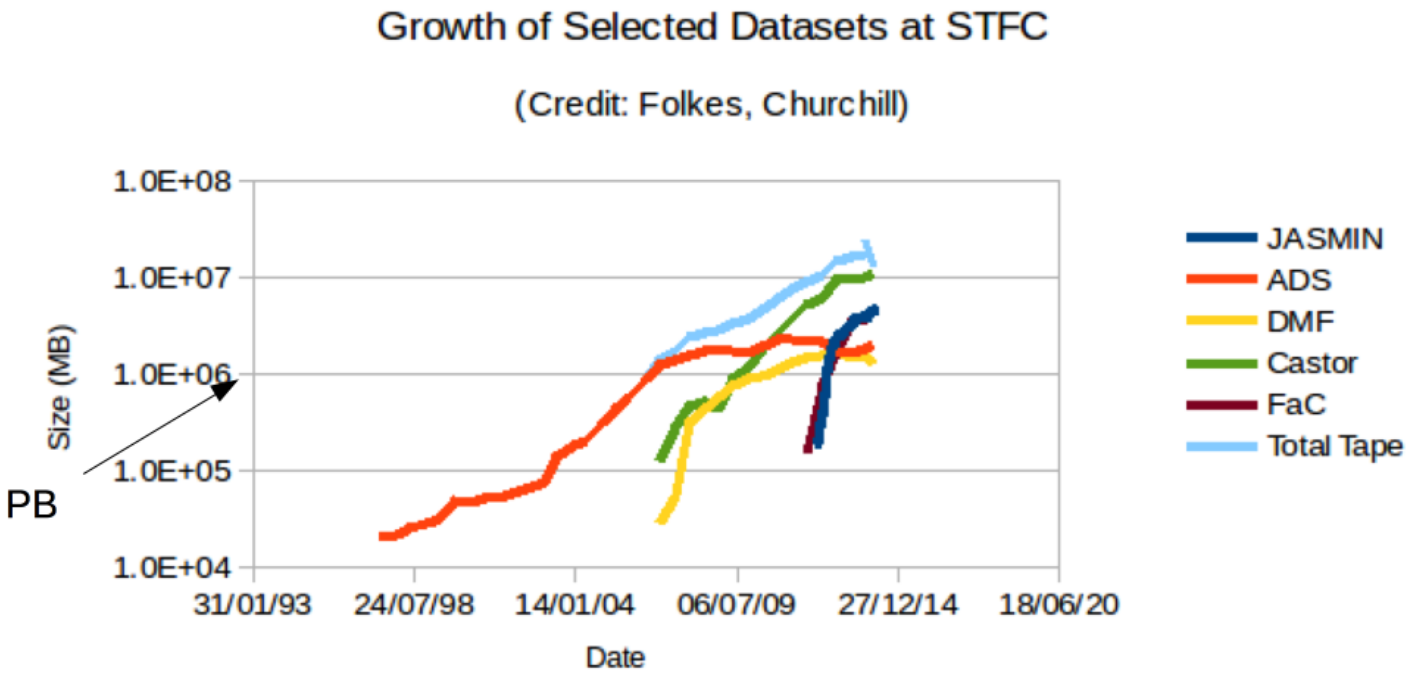

Predictions for JASMIN in 2020? 30 — 85 PB of unique data[1]!
But we think we could only fit only 30 PB disk in the physical space available[2]!

([1]Not including CMIP6, which might be anything from 30 PB up. [2]Unless we can throw out the CERC Tier1 centre with whom we share!)

Sentinel 1

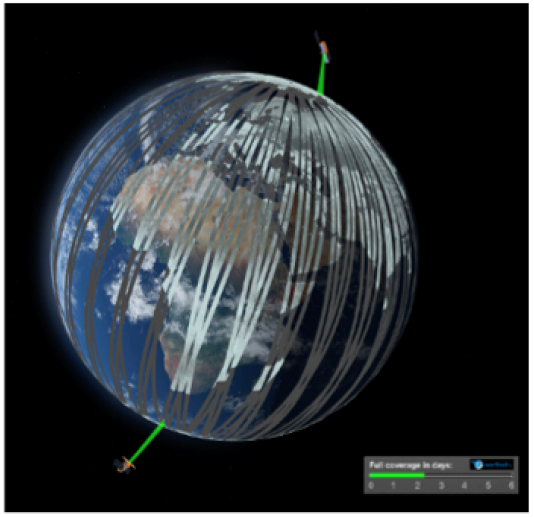

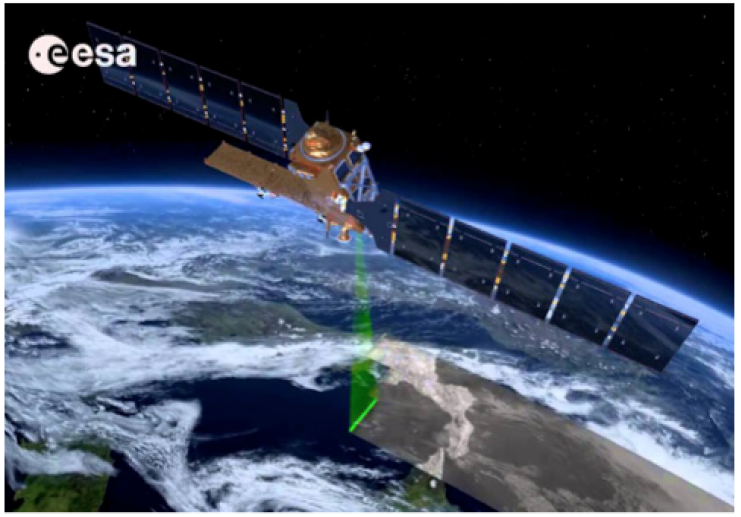

Sentinel 1A: Launched 2014
(1B due 2016)

- Key instrument: Synthetic Aperture Radar
- Data rate (two satellites: raw 1.8 TB/day, archive products ~ 2 PB/year)

NERC SCIENCE OF THE ENVIRONMENT

COMET: Centre for Observation and Modelling of Earthquakes, Volcanoes, and Tectonics

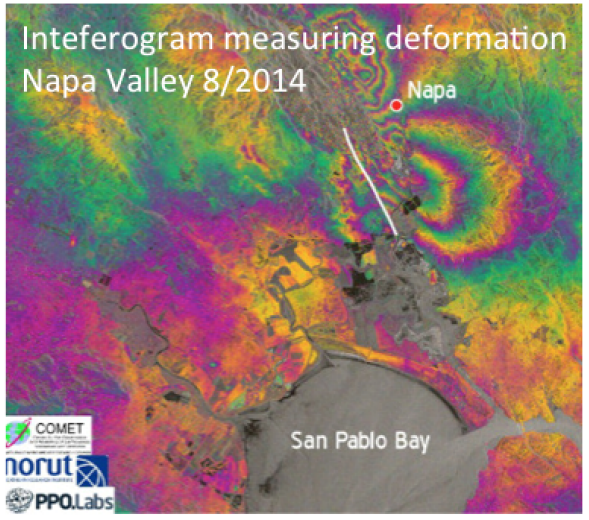

(Picture credits: ESA, Arianespace.com, PPO.labs-Norut–COMET-SEOM Insarap study, ewf.nerc.ac.uk/2014/09/02/new-satellite-maps-out-napa-valley-earthquake/)

Sentinel Data Rates

Satellite	Launch Dates	Daily Data Rate	Product Archive
S1A, S1B	Apr 2014	1.8 TB/day raw	2 PB/year
S2A, S2B	Jun 2015	1.6 TB/day raw	2.4 PB/year
S3A, S3B	Oct 2015	0.6 TB/day raw	2 PB/year (L1,L2,L3)

with more satellites in the pipeline. Too easy to say "petabytes"!

Sentinel Data Rates

Satellite	Launch Dates	Daily Data Rate	Product Archive
S1A, S1B	Apr 2014	1.8 TB/day raw	2 PB/year
S2A, S2B	Jun 2015	1.6 TB/day raw	2.4 PB/year
S3A, S3B	Oct 2015	0.6 TB/day raw	2 PB/year (L1,L2,L3)

with more satellites in the pipeline. Too easy to say "petabytes"!

- Traditional approach: write data to tapestore, users retrieve scenes from a catalogue!
- Modern "big data" aproach: users want to do "whole mission" reprocessing!
  - e.g. QA4ECV (J-P Muller): bought 800 TB of disk in the JASMIN system, now running whole mission reprocessing 100x faster than their in-house cluster. Days to test new science instead of months. Massive improvement in scientific throughput!

U.S. National Academy

"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

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

So we have built an "HPC-data" cloud: JASMIN

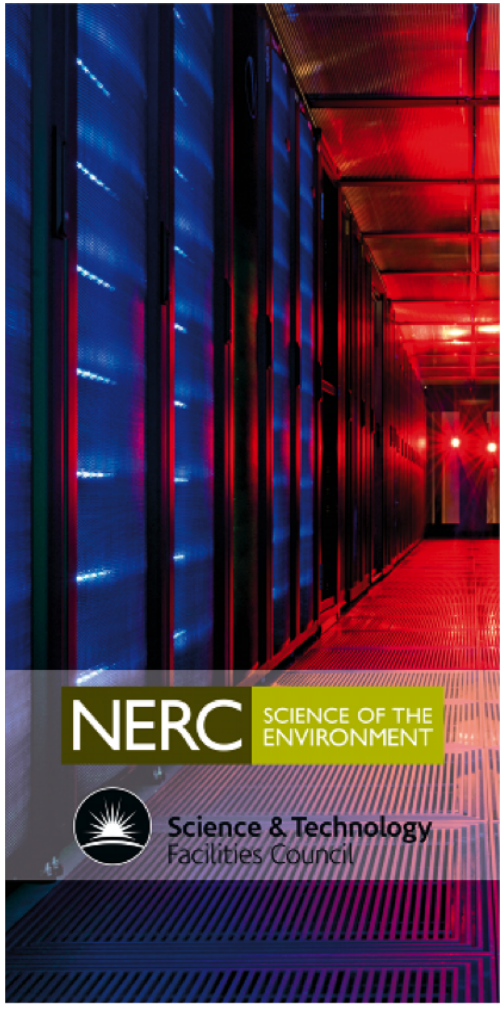

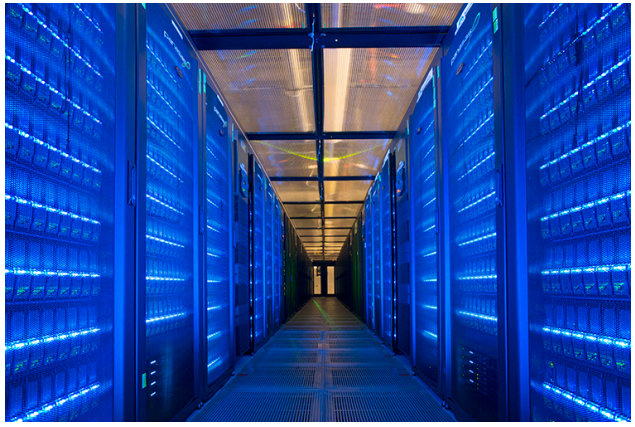

- 16 PB Fast Storage (Panasas, many Tbit/s bandwidth)
- 1 PB Bulk Storage
- Elastic Tape
- 4000 cores: half deployed as hypervisors, half as the "Lotus" batch cluster.
- Some high memory nodes, a range, bottom heavy.

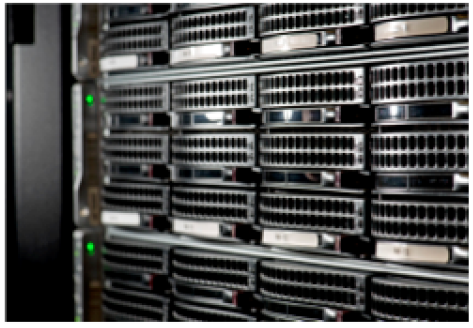

Virtual Organisations

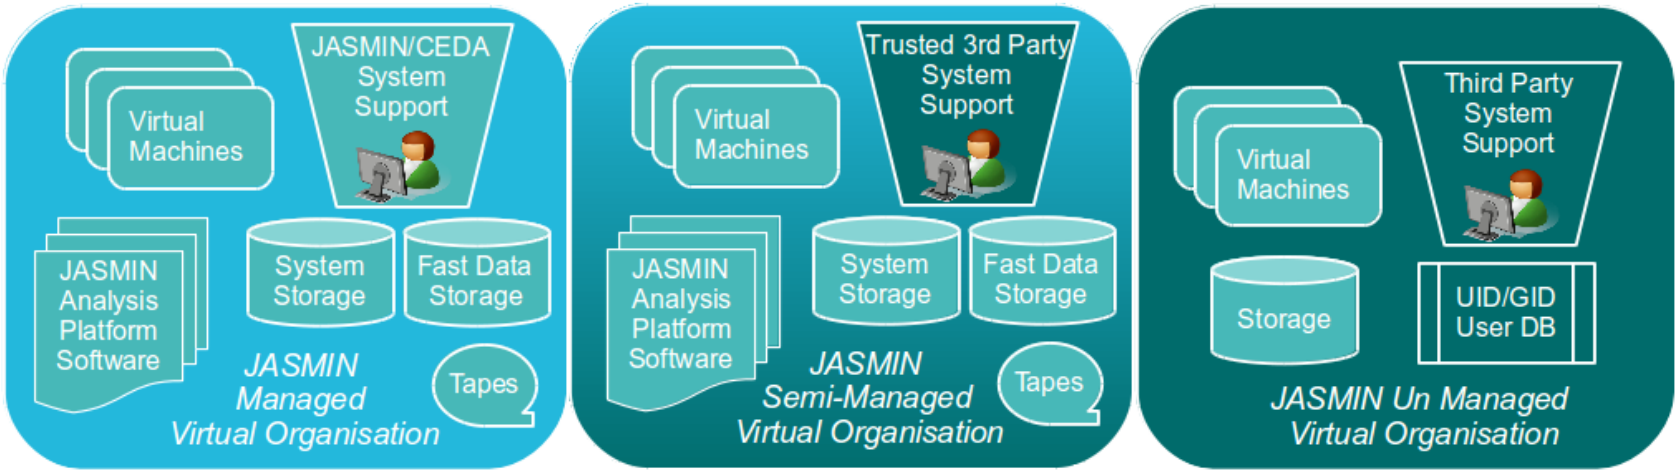

Platform as a Service $\longrightarrow$ Infrastructure as a Service

Example: NCAS will run a semi-managed virtual organisation (with multiple group work spaces), but large groups within NCAS can themselves also run virtual organisations.

High performance, curation + facilitation

Objective is to provide an environment with high performance access to curated data archive **and** a high performance data analysis environment!

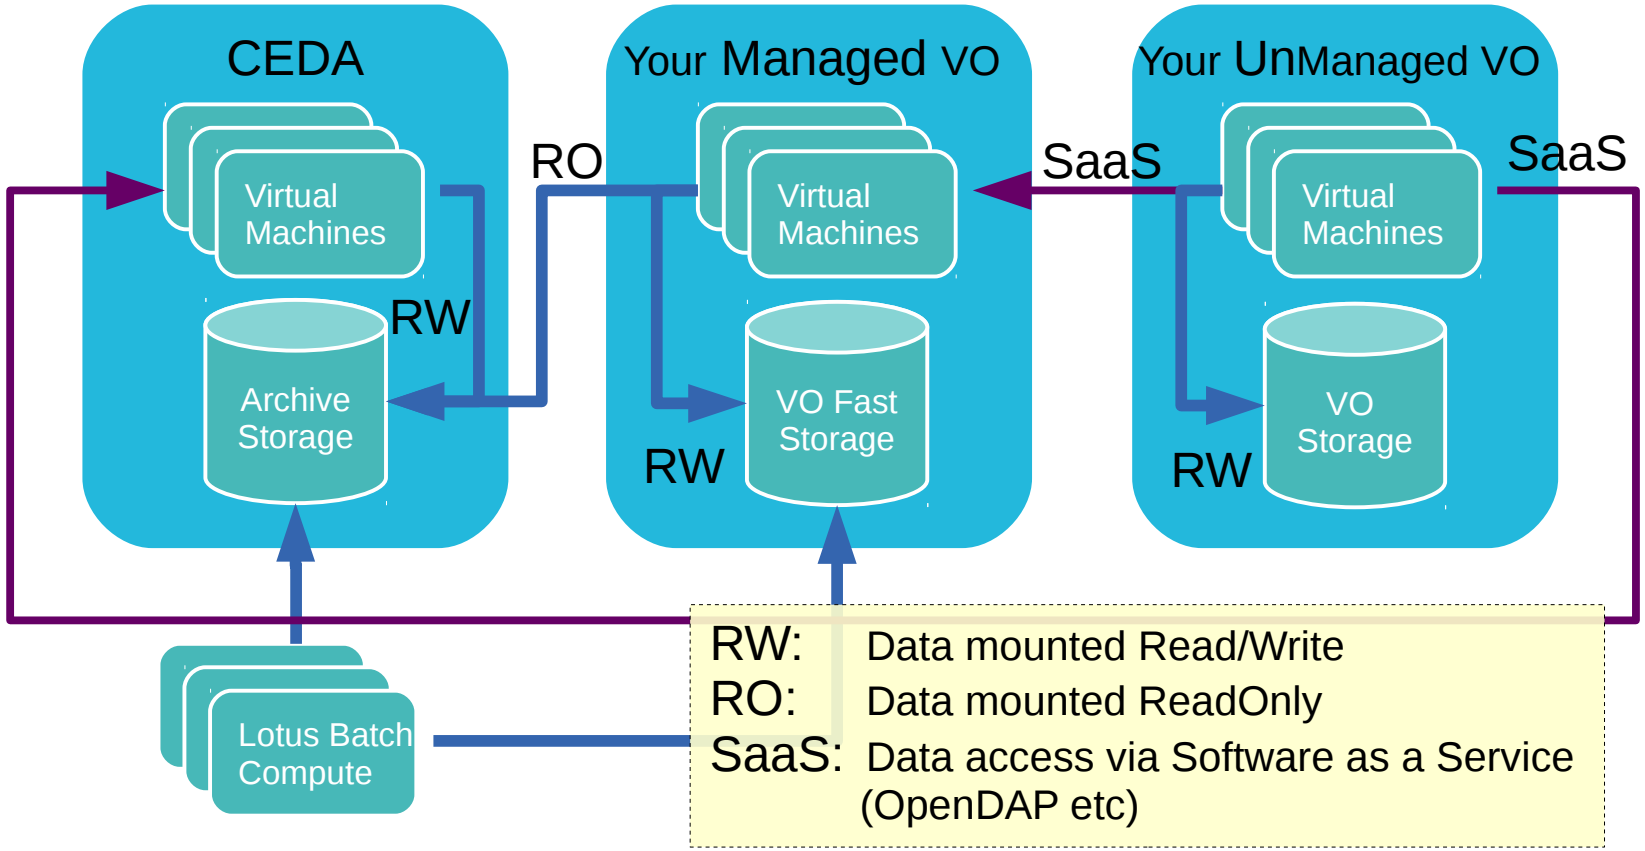

Integrated Cloud Provisioning

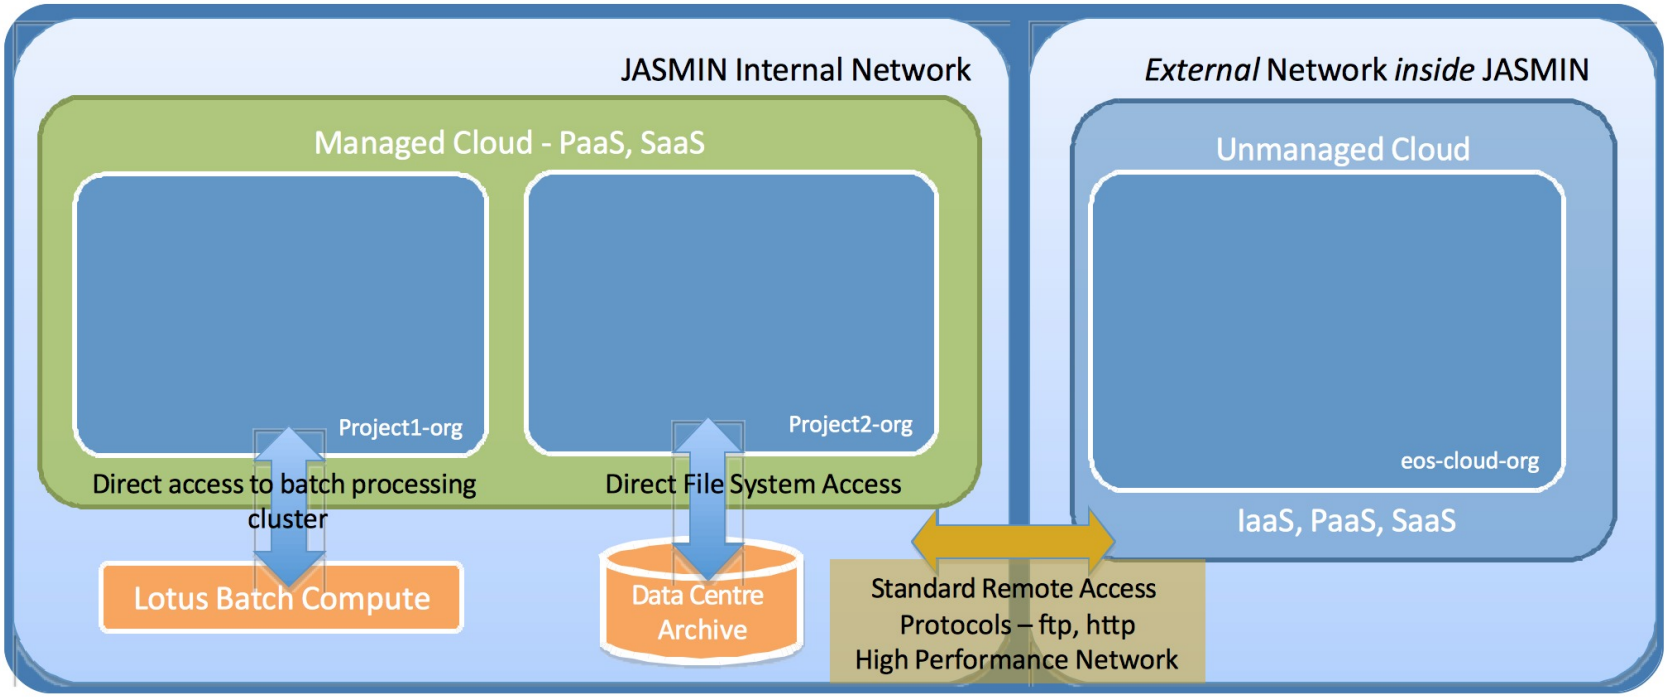

Currently o(100) "Group Work Spaces" in the managed cloud serving o(100) "virtual organisations" and o(500) users (there is some overlap). Unmanaged cloud is currently in testing with a few brave souls.

Structural View

CEDA AS (once BADC)

CEDA EO (once NEODC)

CEDA Solar (once UKSSDC)

IPCC DDC

etc

NERC Managed Analysis Computing

(CEMS + Shared Systems for NCAS, MetO, NOC etc)

NERC Cloud Analysis Computing

(EOS Cloud, Env WB etc)

etc

CEDA Archive Services

Data Centres, Curation, DB systems

User management, External Helpdesk

CEDA Compute Services

Compute Cloud:

PaaS (JAP +Generic Science VMs + User Management), IaaS

External Helpdesk

JASMIN Compute and Storage

Lotus + Private Cloud + Tape Store + DMZ for data transfer

Internal Helpdesk

NERC SCIENCE OF THE ENVIRONMENT

Big Cache

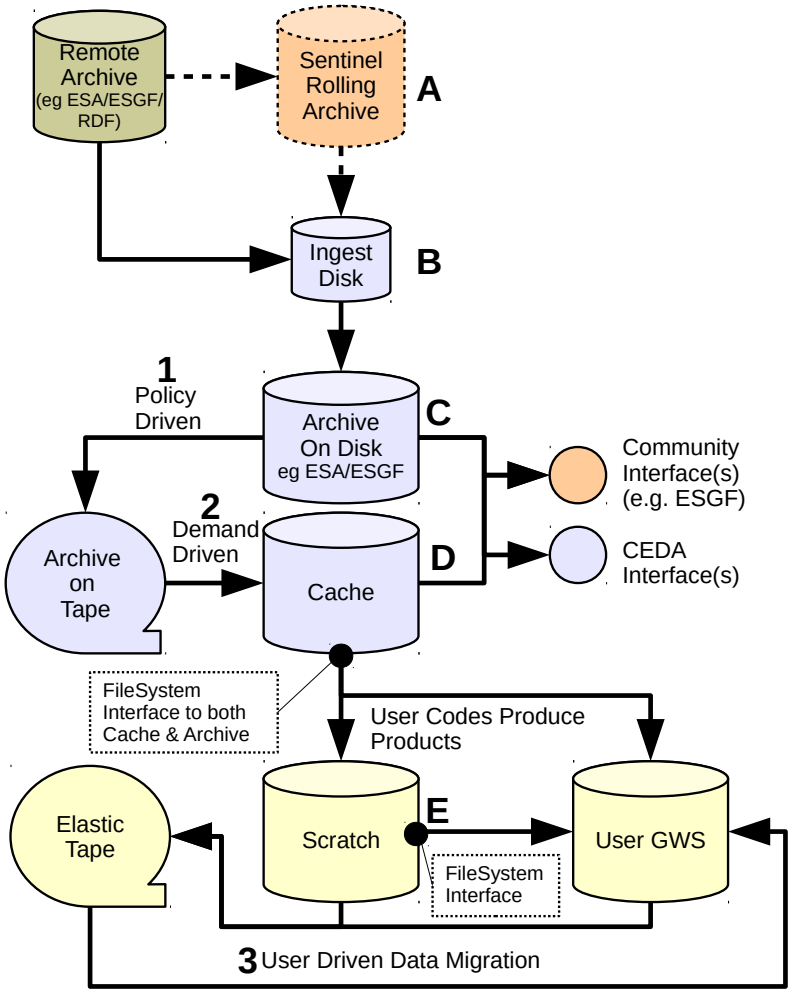

We need to build this out in the next six months, before we end up with PB and unhappy users!

Cloudbursting

Final Remarks

- When we consider the entire workflow associated with environmental science, we realise that the "time in the supercomputer" **doing** simulation, or "the time in the satellite ground segment" or the "time at sea collecting data" is only a small part of the entire workflow.
- When we look at the trend in the balance of hardware spending nearly anywhere people are analysing (or producing) data, we see a trend towards a greater proportion of the funding on the storage, but
- We have yet to see a commensurate trend towards the spend for an appropriate software infrastructure for data, and
- We have yet to see an understanding by the institutions of the costs ahead, yet alone their already sunk costs.
- When the institutions get on top of this, they will come running "cloud-ward", but it's important to make sure that cloud is "data-intensive-ready"!