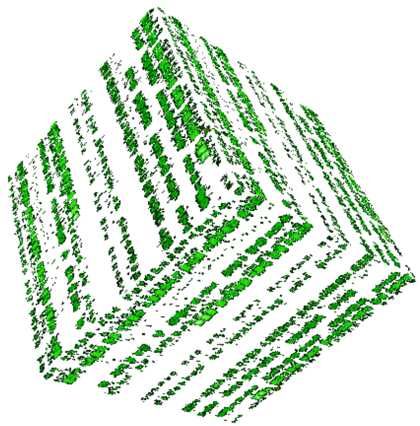

[gamingfeeds.com]

Science SQL:
Advancing from Data to Service Stewardship

LSDMA Symposium "The Challenge of Big Data in Science", Karlsruhe, 2015-oct-01

Peter Baumann

Jacobs University | rasdaman GmbH

baumann@rasdaman.com

[co-funded by EU through EarthServer 1/2, PublicaMundi]

Array Analytics Research @ Jacobs U

- Large-Scale Scientific Information Systems research group
  - Flexible, scalable n-D array services
  - www.jacobs-university.de/lsis
- Main results:
  - pioneer Array DBMS, rasdaman
  - standardization:
    - OGC Big Geo Data (also ISO, INSPIRE, W3C)
    - ISO „Science SQL“

Hiring PhD students, PostDocs

Big Data, from a Database Perspective

Structural Variety in Big Data

- Stock trading: 1-D sequences (i.e., arrays)
- Social networks: large, homogeneous graphs
- Ontologies: small, heterogeneous graphs
- Climate modelling: 4D/5D arrays
- Satellite imagery: 2D/3D arrays (+irregularity)
- Genome: long string arrays
- Particle physics: sets of events
- Bio taxonomies: hierarchies (such as XML)
- Documents: key/value stores = sets of unique identifiers + whatever
- etc.

Structural Variety in Big Data

- Stock trading: 1-D sequences (i.e., arrays)
- Social networks: large, homogeneous graphs
- Ontologies: small, heterogeneous graphs
- Climate modelling: 4D/5D arrays
- Satellite imagery: 2D/3D arrays (+irregularity)
- Genome: long string arrays
- Particle physics: sets of events
- Bio taxonomies: hierarchies (such as XML)
- Documents: key/value stores = sets of unique identifiers + whatever
- etc.

Array Services

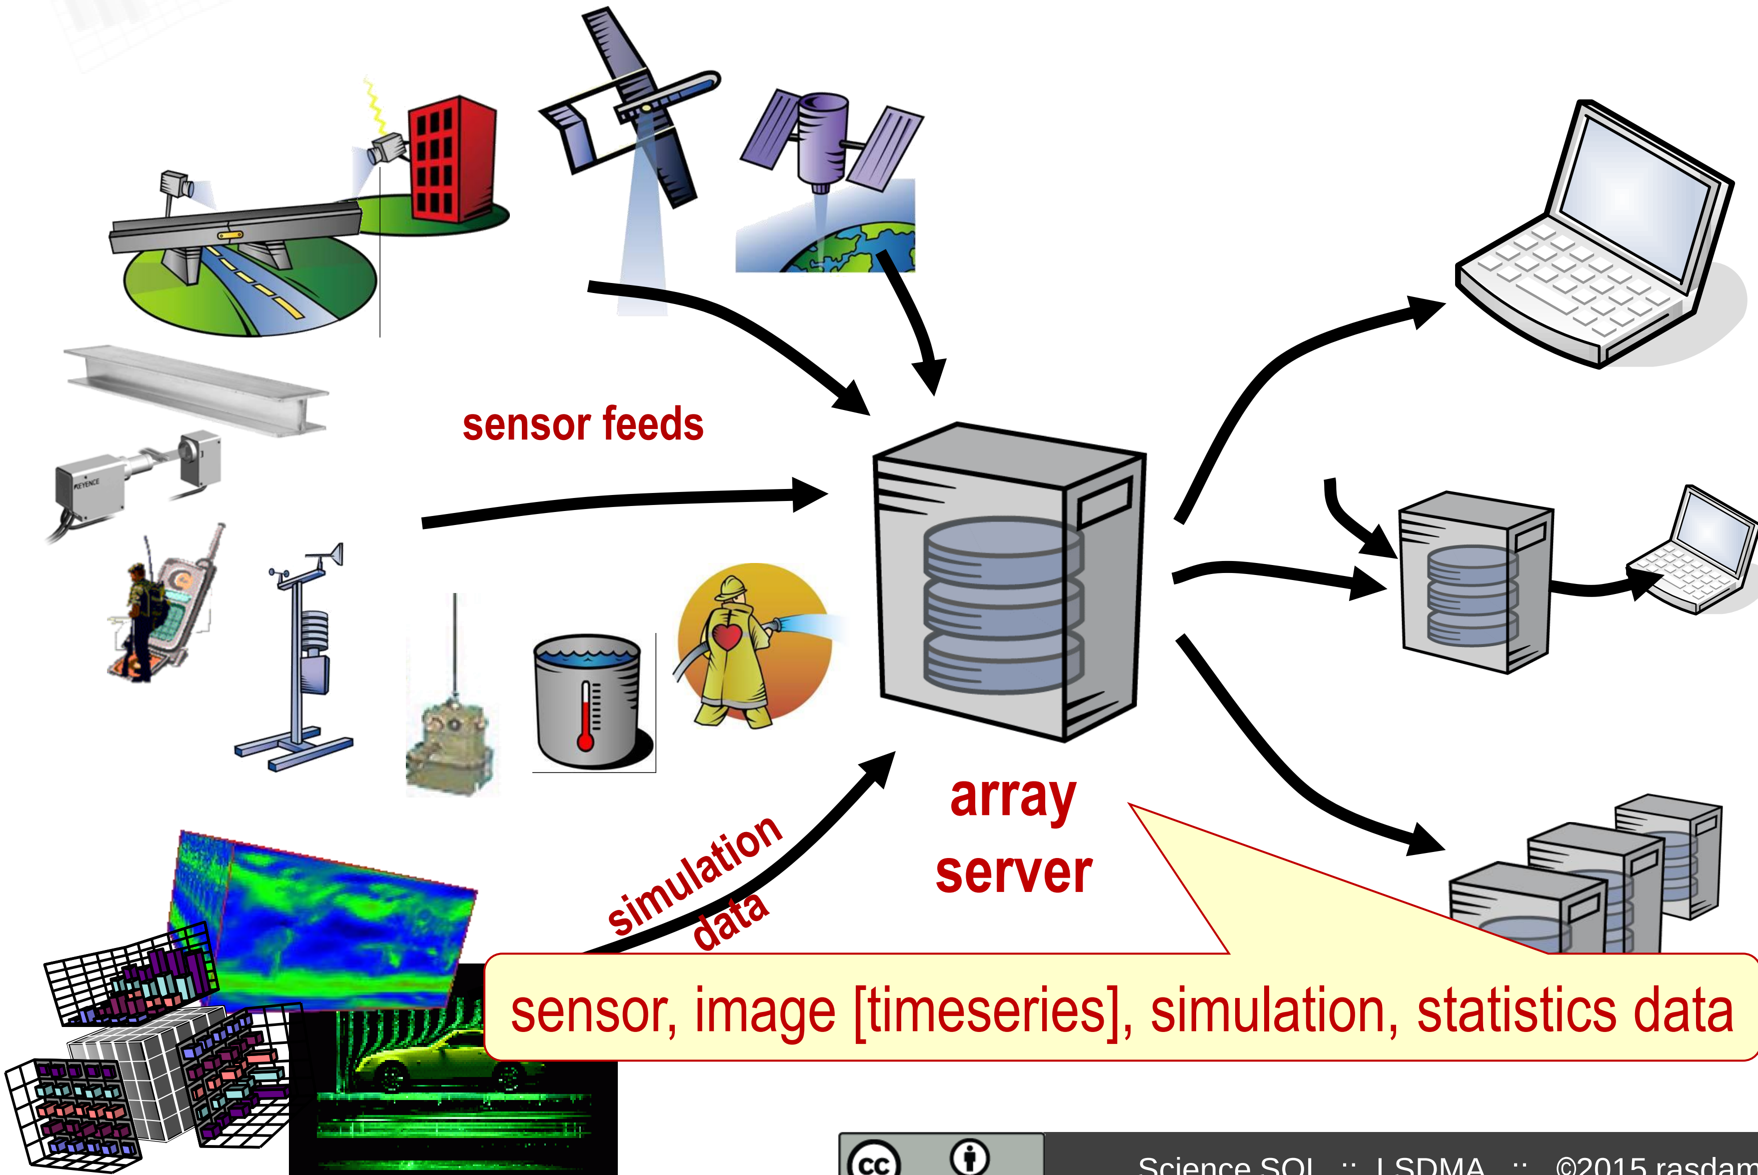

Counter Example

- OGC WCPS: semantics in query → machine understandable

```
for $c in ( M1, M2, M3 )
return encode abs( $c.red - $c.nir ), "hdf" )
```

- OGC WPS: semantics in human-readable text

```
<ProcessDescriptions ...>
  <ProcessDescription processVersion="2" storeSupported="true" statusSupported="false">
    <ows:Identifier>Buffer</ows:Identifier>
    <ows:Title>Create a buffer around a polygon.</ows:Title>
    <ows:Abstract>Create a buffer around a single polygon. Accepts the polygon as GML and provides GML output for the buffered feature. </ows:Abstract>
    <ows:Metadata xlink:title="spatial" />
    <ows:Metadata xlink:title="geometry" />
    <ows:Metadata xlink:title="buffer" />
    <ows:Metadata xlink:title="GML" />
    <DataInputs>
      <Input>
        <ows:Identifier>InputPolygon</ows:Identifier>
        <ows:Title>Polygon to be buffered</ows:Title>
        <ows:Abstract>URI to a set of GML that describes the polygon.</ows:Abstract>
        <ComplexData defaultFormat="text/XML" defaultEncoding="base64" defaultSchema="http://foo.bar/gml/3.1.0/polygon.xsd">
          <SupportedComplexData>
                                                                       1,1           Top
```

Inset: Hadoop not the Answer to All

- no builtin knowledge about structured data types
  - "Since it was not originally designed to leverage the structure […] its performance […] is therefore suboptimal" [Daniel Abadi]
  - M. Stonebraker (XLDB 2012): „will hit a scalability wall"

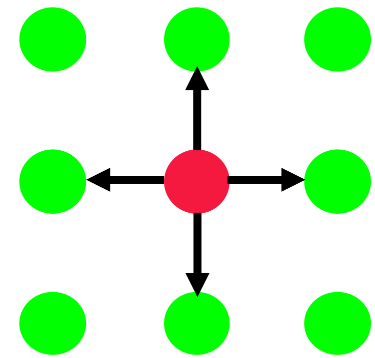

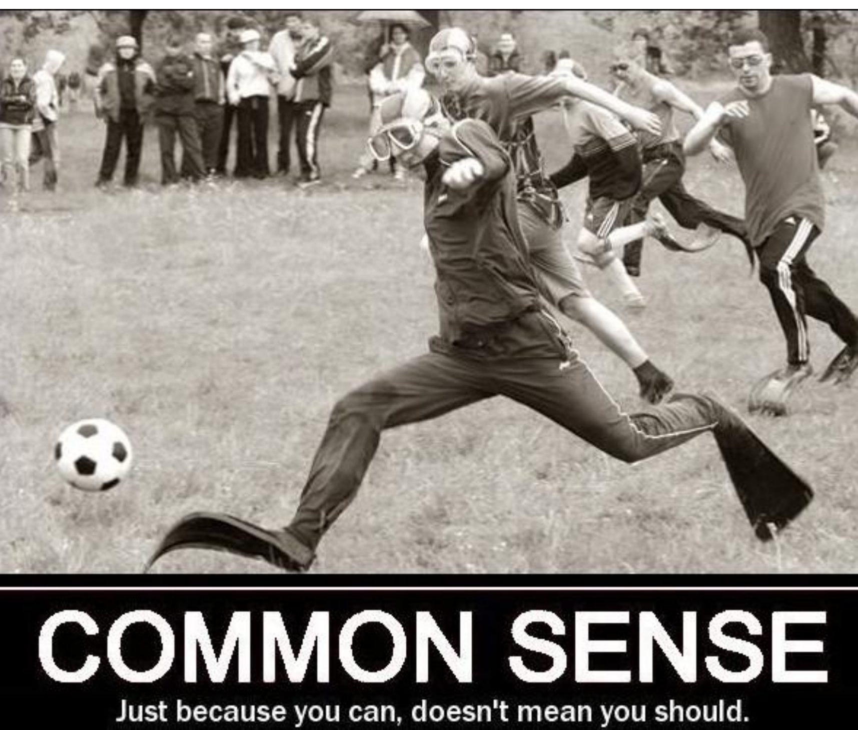

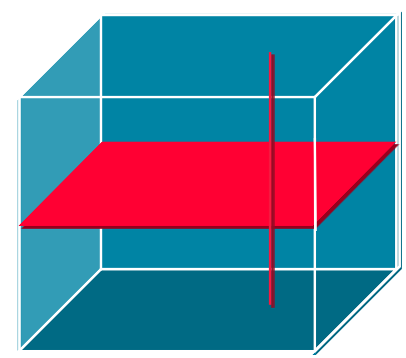

The rasdaman Array Database

Agile Array Analytics: rasdaman

- „raster data manager“: SQL + n-D arrays
- Scalable parallel “tile streaming” architecture
- Integrates with R, python, …
- Blueprint for ISO Array SQL standard

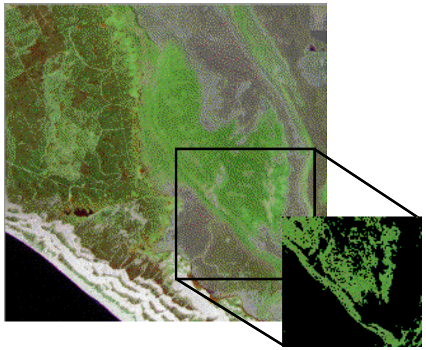

Visualization-as-a-Query

```
for $s in (SatImage), $d in (DEM)
where $s/metadata/@region = "Glasgow"
return
   encode(
      struct {
         red:    (char) $s.b7[x0:x1,x0:x1],
         green:  (char) $s.b5[x0:x1,x0:x1],
         blue:   (char) $s.b0[x0:x1,x0:x1],
         alpha:  (char) scale( $d, 20 )
      },
      "image/png"
   )
```

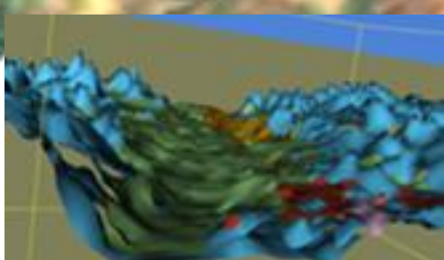

[JacobsU, Fraunhofer; data courtesy BGS, ESA]

„Science SQL“ in ISO
[SSDBM 2014]

ISO/IEC JTC 1/SC 32

Date: 2014-06-04

WD 9075-15:2014(E)

ISO/IEC JTC 1/SC 32/WG 3

The United States of America (ANSI)

Information technology — Database languages — SQL —

Part 15:
Multi-Dimensional Arrays (SQL/MDA)

Technologies de l'information — Langages de base de données — SQL —

Partie 15: Tableaux multi-dimensionnels (SQL/MDA)

```
create table LandsatScenes(
    id: integer not null, acquired: date,
    scene: row( band1: integer, ..., band7: integer ) mdarray [ 0:4999,0:4999] )
```

```
select  id, encode(scene.band1-scene.band2)/(scene.nband1+scene.band2)), „image/tiff“ )
from    LandsatScenes
where acquired between „1990-06-01“ and „1990-06-30“ and
          avg( scene.band3-scene.band4)/(scene.band3+scene.band4)) > 0
```

Integrating External Codes

- User-Defined Functions = external code dynamically linked into server
  - rasdaman: Same API as clients, auto-generated adapter code → easy to use
  - Integrated with tile management, parallelization, ...
- Ex: “*NDVI from raw Landsat subset, orthorectified with Orfeo Toolbox*“

```
select
  encode(
    otb.orthoRectifFilter(
      ((img.red-img.nir)/(img.red+img.nir))[x0:x1,y0:y1],
      outputSpacing, deformationFieldSpacing
    ),
    „image/png"
  )
from   LandsatRawArchive
```

Tiling: Tuning Data for Applications

- tiling strategies as service tuning [Furtado]:
  - regular directional area of interest

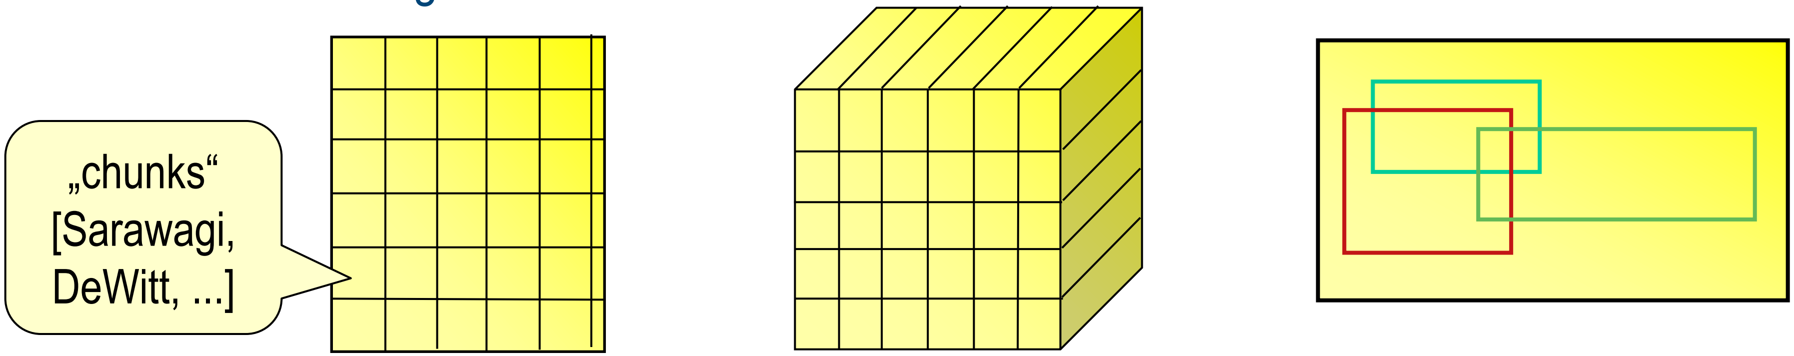

- rasdaman storage layout language

```
insert into MyCollection
  values ...
  tiling area of interest [0:20,0:40], [45:80,80:85]
  tile size 1000000
  index d_index storage array compression zlib
```

Parallel / Distributed Query Processing

```
select
    max((A.nir - A.red) / (A.nir + A.red))
  - max((B.nir - B.red) / (B.nir + B.red))
  - max((C.nir - C.red) / (C.nir + C.red))
  - max((D.nir - D.red) / (D.nir + D.red))
from A, B, C, D
```

1 query → 1,000+ cloud nodes;
CPU + modern hw

[SIGMOD DANAC 2014]

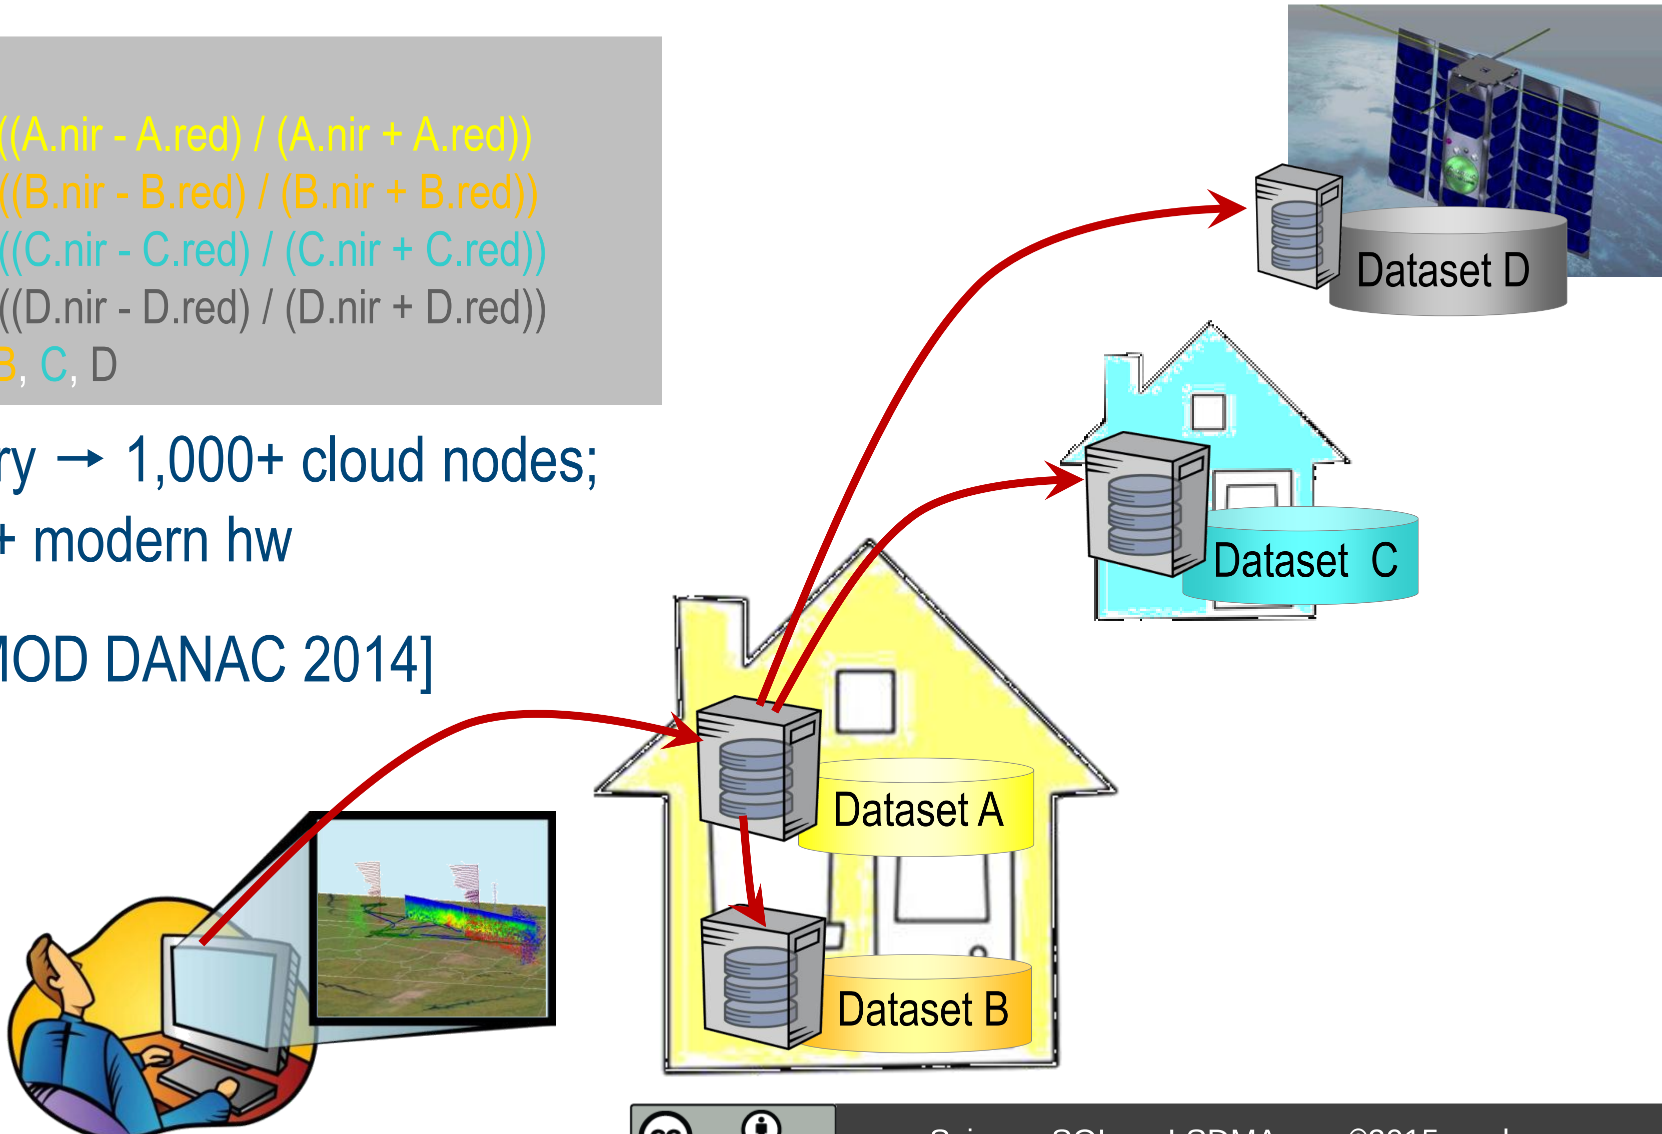

Secured Archive Integration

First-ever direct, ad-hoc mix from protected NASA & ESA services in OGC WCS/WCPS Web client (EarthServer + CobWeb)

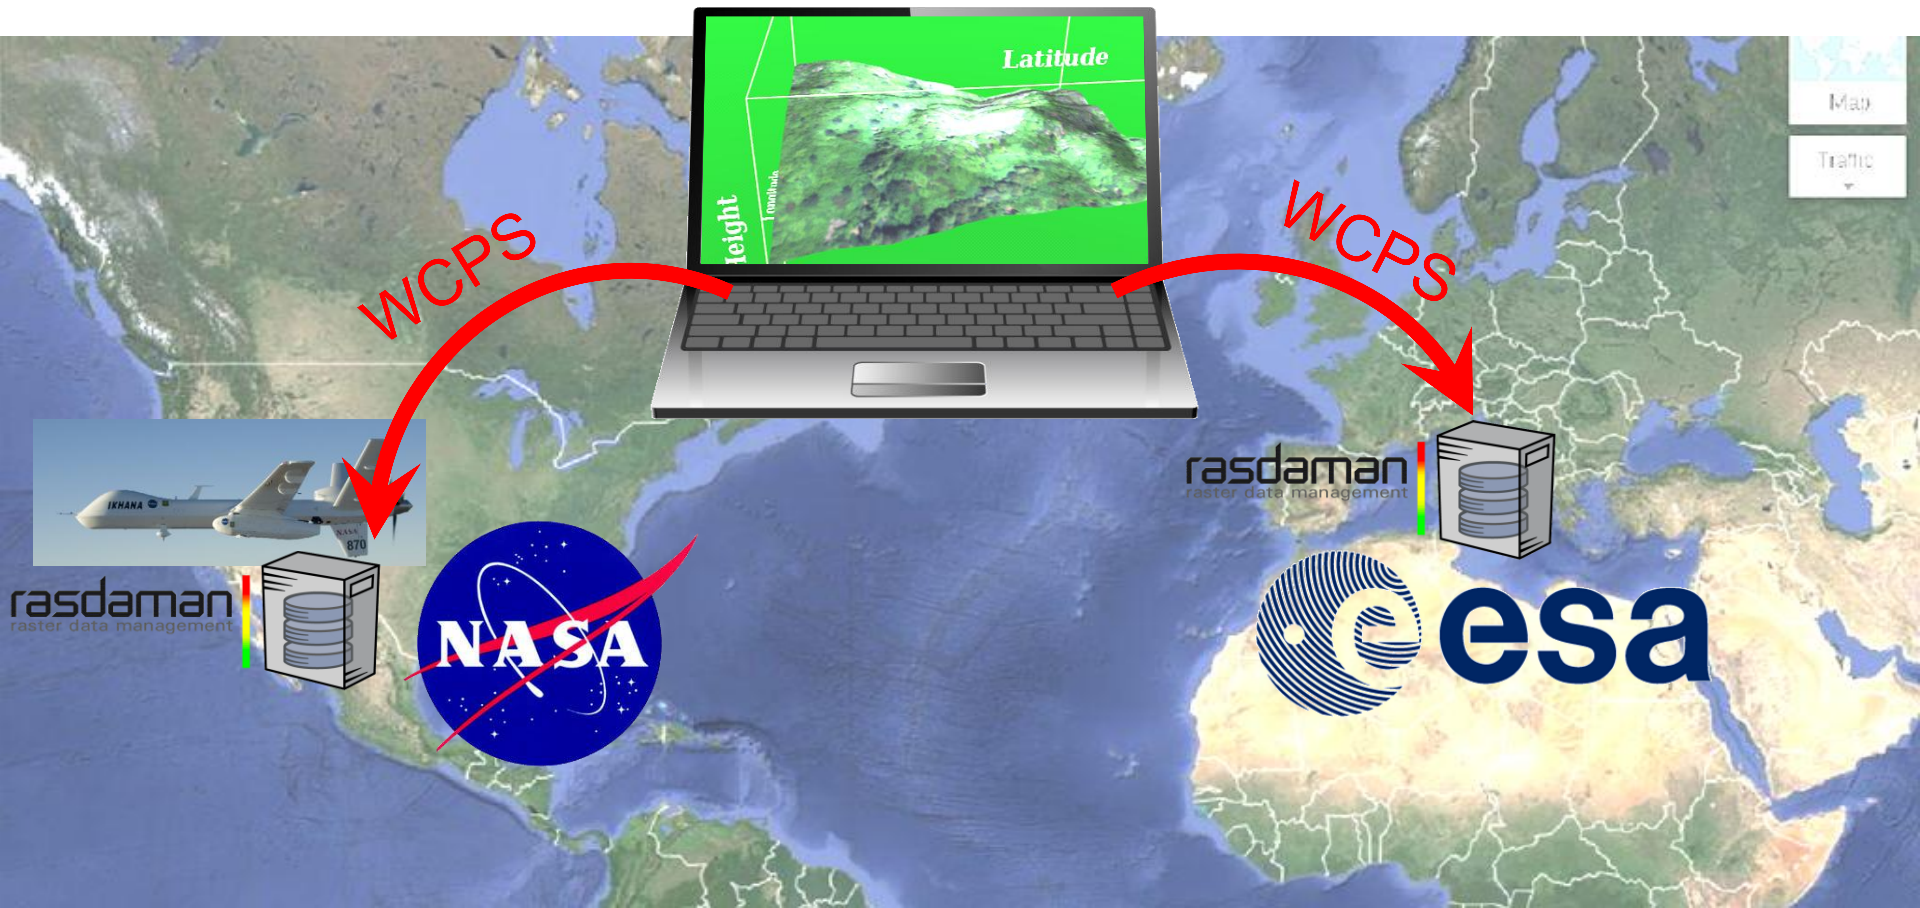

Scalable Geo Service Architecture

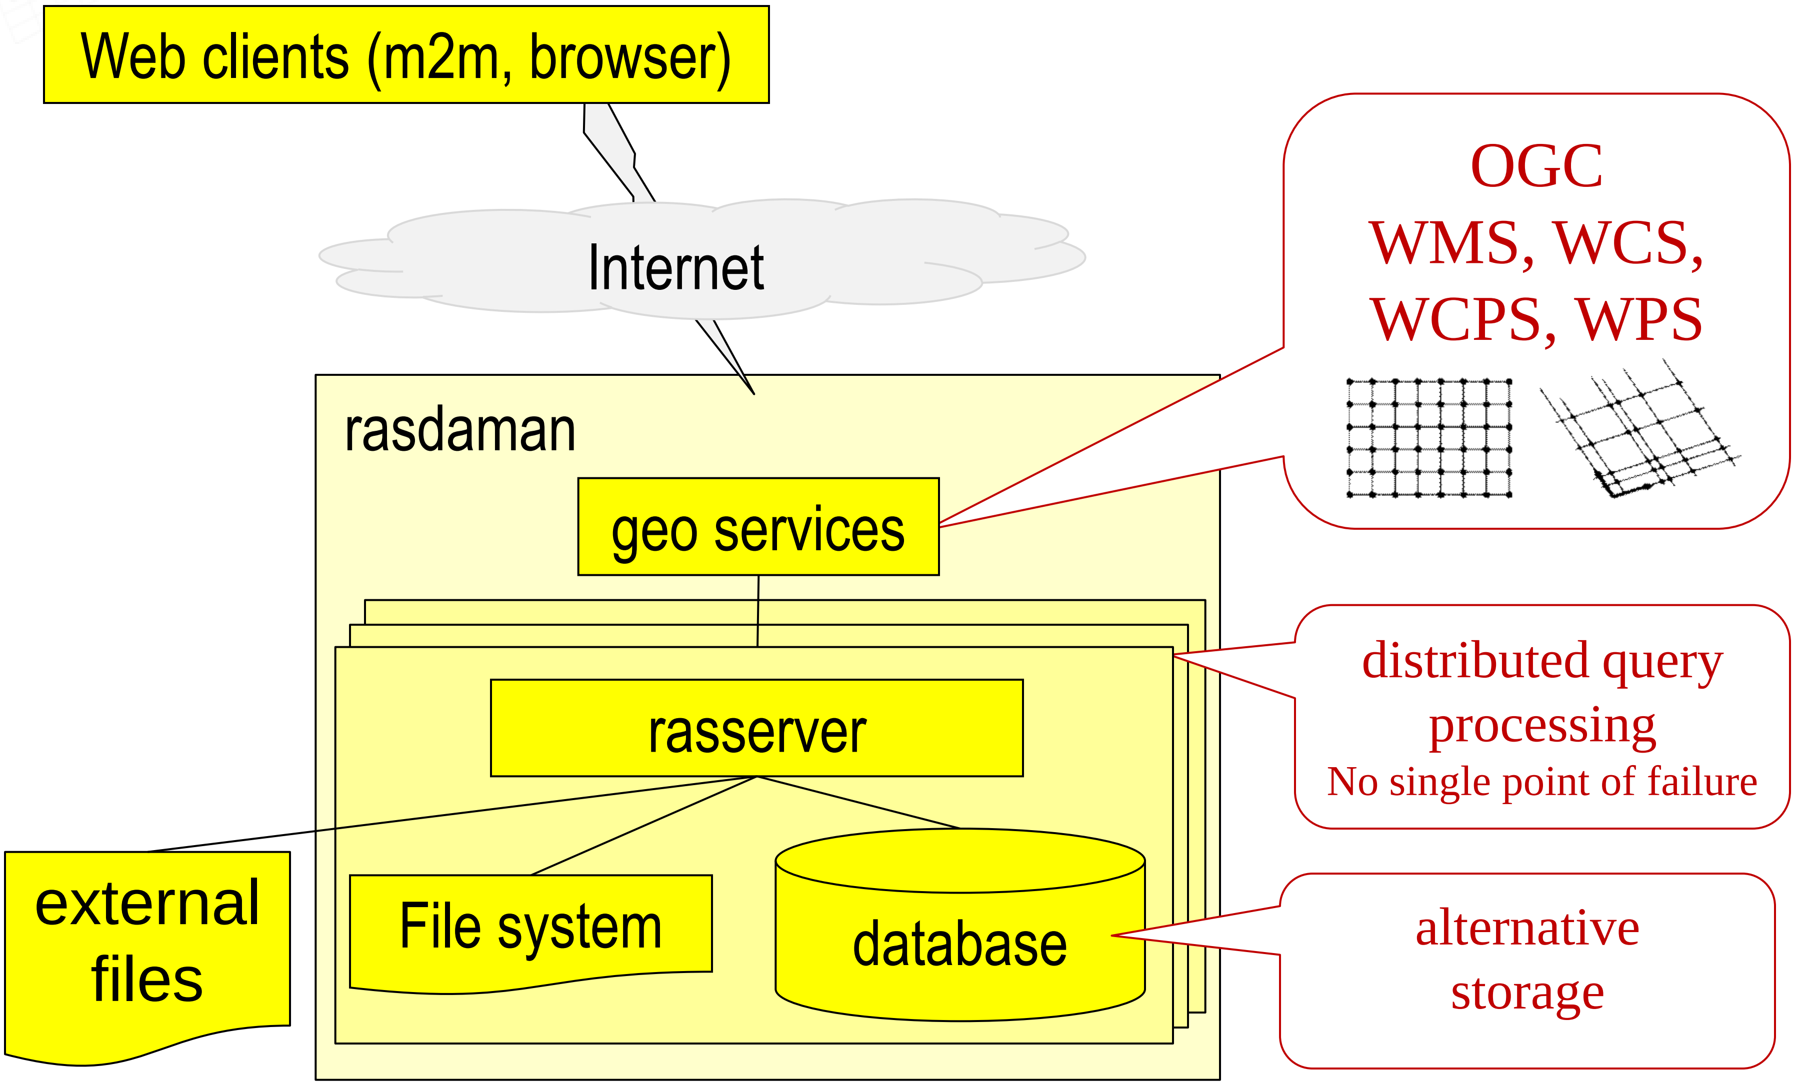

Science Data Use Cases

EarthServer: Datacubes At Your Fingertips

- Agile Analytics on Earth & Planetary datacubes
  - rasdaman + NASA WorldWind
  - Rigorously standards: OGC WMS + WCS + WCPS
  - 100s of TB online now, goal: 1+ Petabyte per cube

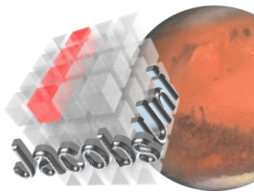

- Intercontinental initiative, 3+3 years: EU + US + AUS

Phase 1 review:
"proven evidence" that rasdaman will "*significantly transform [how to] access and use data*" …and *"with no doubt has been shaping the Big Earth Data landscape"* …

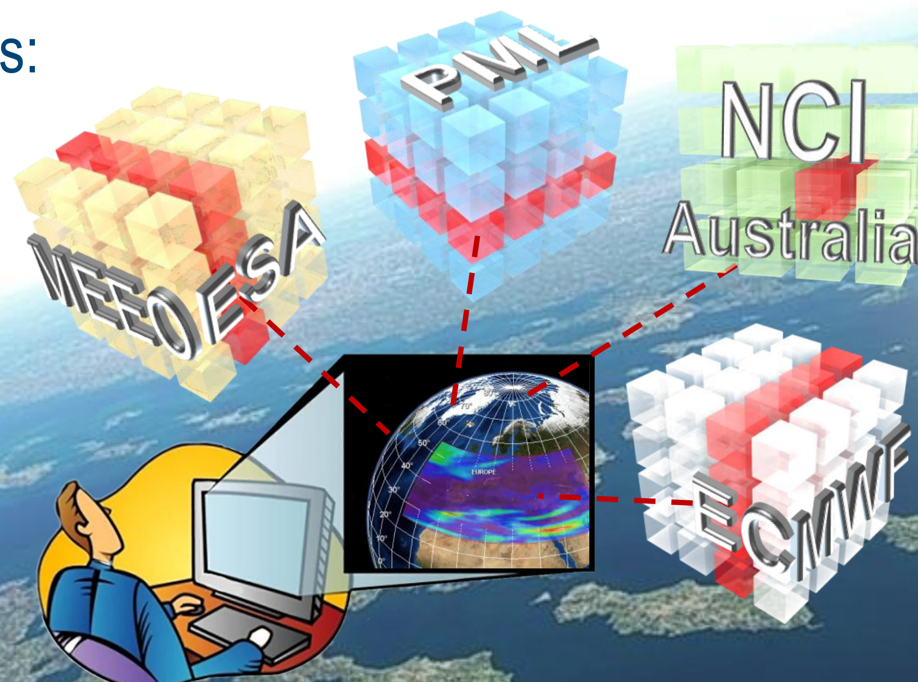

EarthServer Phase 1 & 2 Partners

rasdaman raster data management
JACOBS UNIVERSITY

JACOBS UNIVERSITY

rasdaman raster data manager

MEEO Meteorological Environmental Earth Observation

NASA

PML | Plymouth Marine Laboratory

NCI

CITE

ECMWF

Ερευνητικό Κέντρο Αθηνά
Athena Research Center

BGS 1835 British Geological Survey NATURAL ENVIRONMENT RESEARCH COUNCIL

Fraunhofer

EOX

CNR-IIA

COMETA

Use Case: Human Brain Imaging

- Research goal: understand structural-functional relations in human brain
- Experiments capture activity patterns (PET, fMRI)
  - Temperature, electrical, oxygen consumption, ...
- Ex: "*a parasagittal view of all scans containing critical Hippocampus activations, TIFF-coded.*"

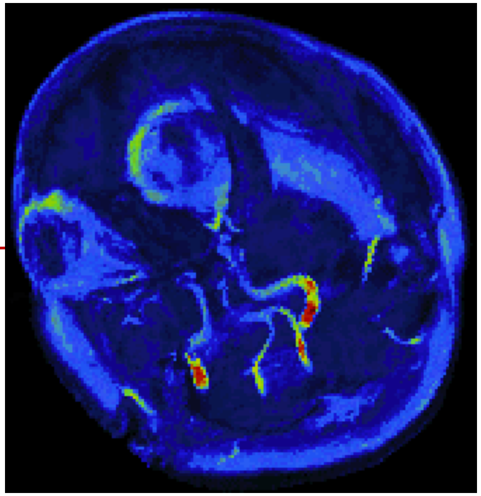

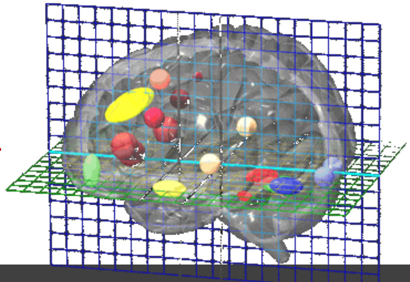

```
select tiff( ht[ $1, *:*, *:* ] )
  from    HeadTomograms as ht,
          Hippocampus as mask
  where   count_cells( ht > $2 and mask )
          / count_cells( mask )
          > $3
```

$1 = slicing position, $2 = intensity threshold value, $3 = confidence

Domains Investigated

- Geo
  - Environmental sensor data, 1-D
  - Satellite / seafloor maps, 2-D
  - Geophysics (3-D x/y/z)
  - Climate modelling (4-D, x/y/z/t)
- Life science
  - Gene expression simulation (3-D)
  - Human brain imaging (3-D / 4-D)
- Other
  - Computational Fluid Dynamics (3-D)
  - Astrophysics (4-D)
  - Statistics (n-D)

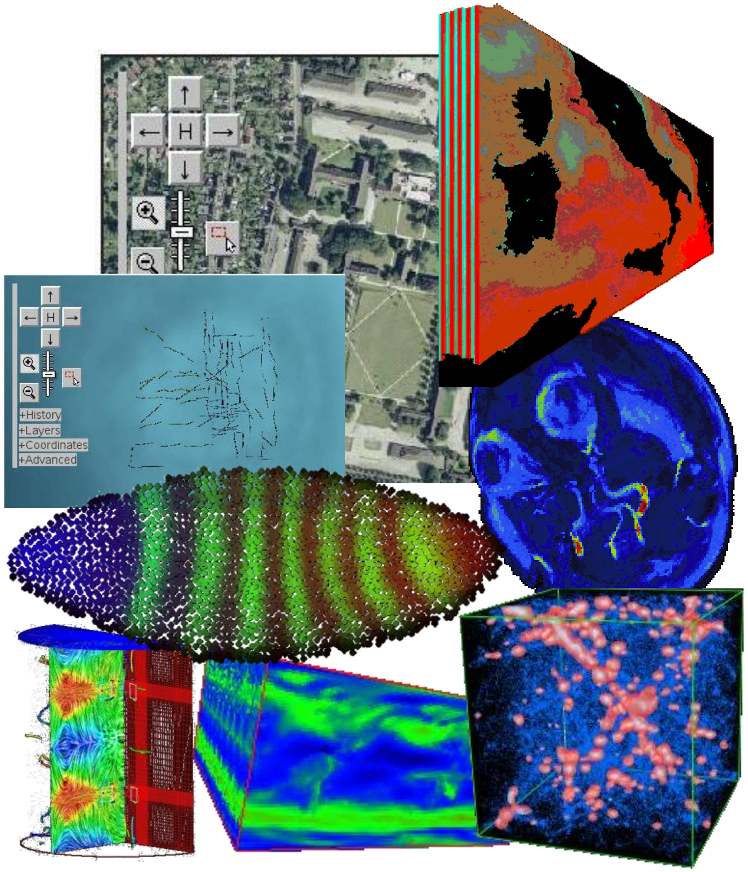

Conclusions

Everbody will have the data - users will go to best service

Data stewardship → service stewardship

flexibility, performance, visualization, tools, ...

cf. Array Databases like rasdaman

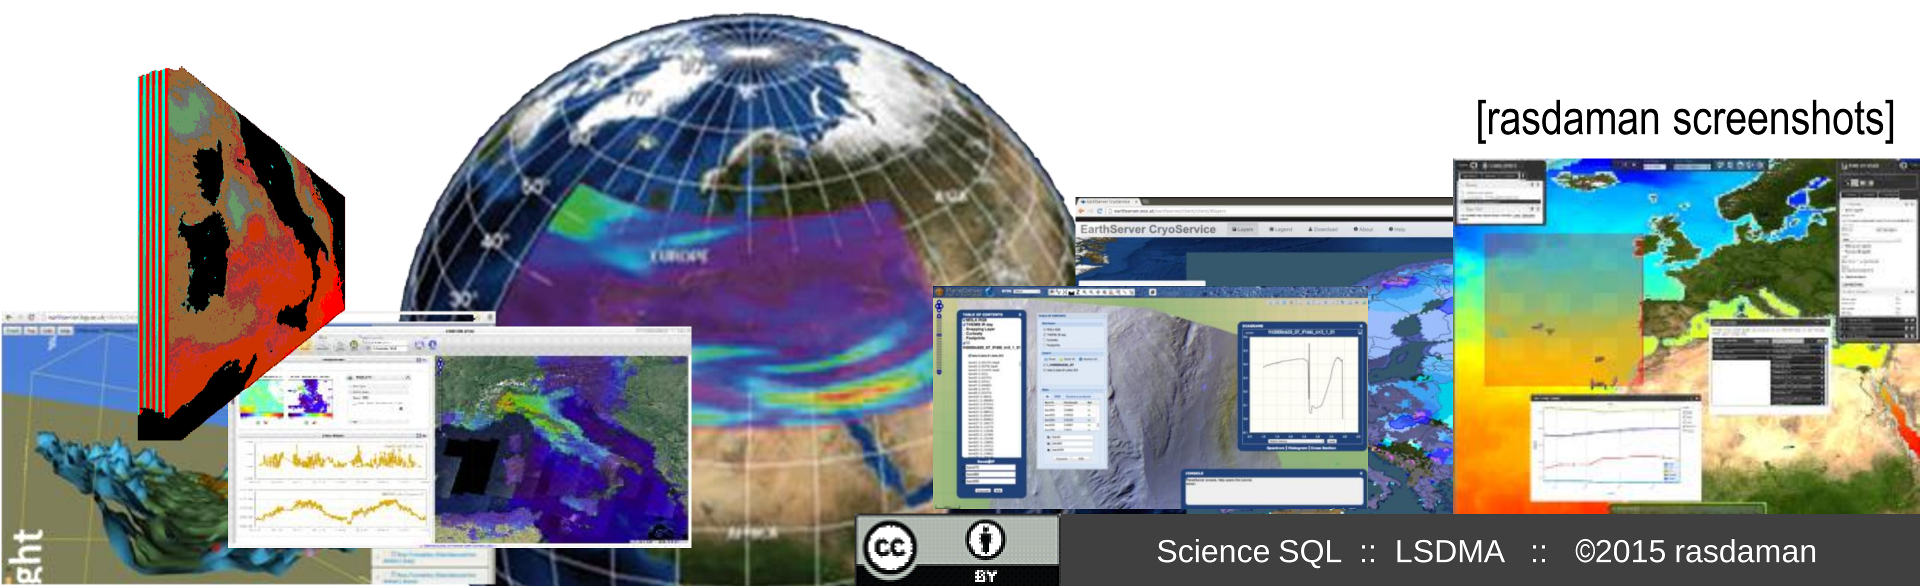

[rasdaman screenshots]