



INDIGO - DataCloud

Ophidia: a big data analytics framework for eScience

Sandro Fiore
(sandro.fiore@cmcc.it)

Advanced Scientific Computing Division
Director

*Euro Mediterranean Center on Climate Change
(CMCC)*

Outline

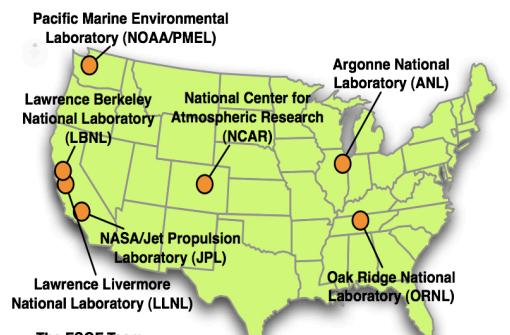
- Background, data analytics requirements & use cases
- Ophidia introduction
- Architecture
- Operators & primitives
- Terminal
- Interoperability and standards
- CMCC activities
- Demo



ENES & IS-ENES/ESGF eco-system

II CMIP5 Federated Archive

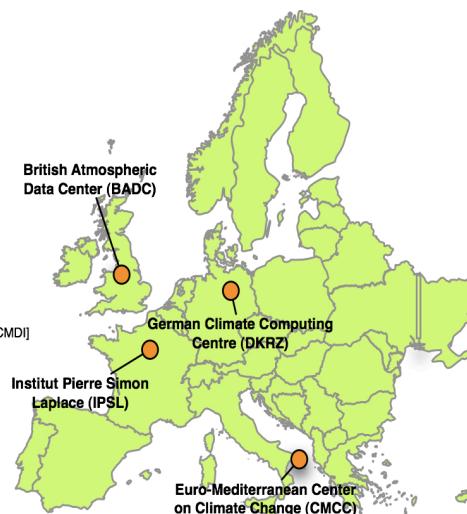
Summary	
Modeling centers	26
Models	51
Experiments	98
Data nodes	20
P2P Index	10
Datasets	44162
Size	1,754.33 TB
Files	3,095,405



The ESGF Team

- Gavin M. Bell, Bob Drach, Charles Doutriaux, Renata McCoy, Dean Williams [LLNL/PCMDI]
- Luca Cincunini, Dan Crichton, Chris Mattmann [NASA/JPL]
- John Harney, Galen Shipman, Feiyi Wang, [ORNL]
- Roland Schweitzer [NOAA/PMEL]
- Rachana Ananthkrishnan, Neill Miller, Ian Foster [ANL]
- Estanti Gonzales, Stephen Kindermann [DKRZ]
- Philip Kershaw, Stephen Pascoe [BADC]
- Sandro Fiore, Giovanni Aloisio [CMCC]
- et al

The ESGF Sponsors



5

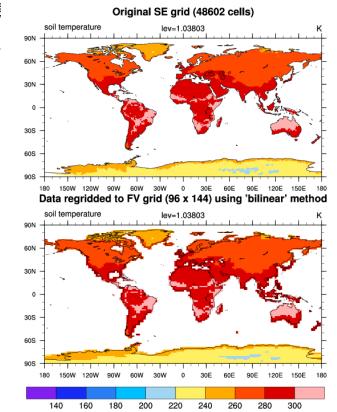
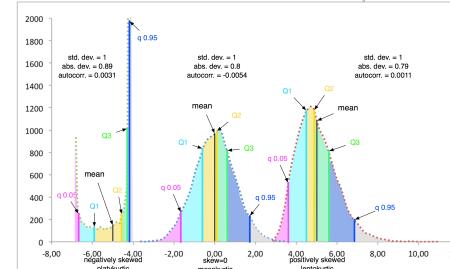
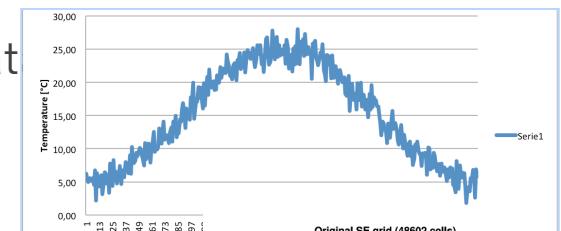
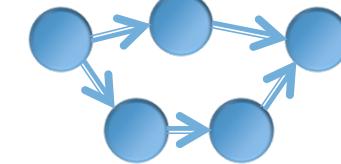
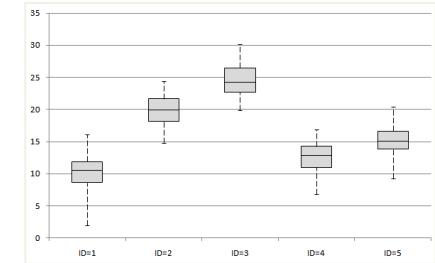
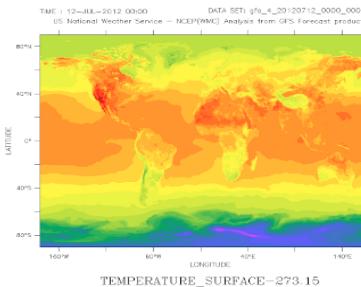
Model	# Datasets	Replica	Size (GB)	# Files
CMCC-CESM	48		694.53	9,942
CMCC-CM	519		91,088.84	130,148
CMCC-CMS	46		1,728.73	8,595





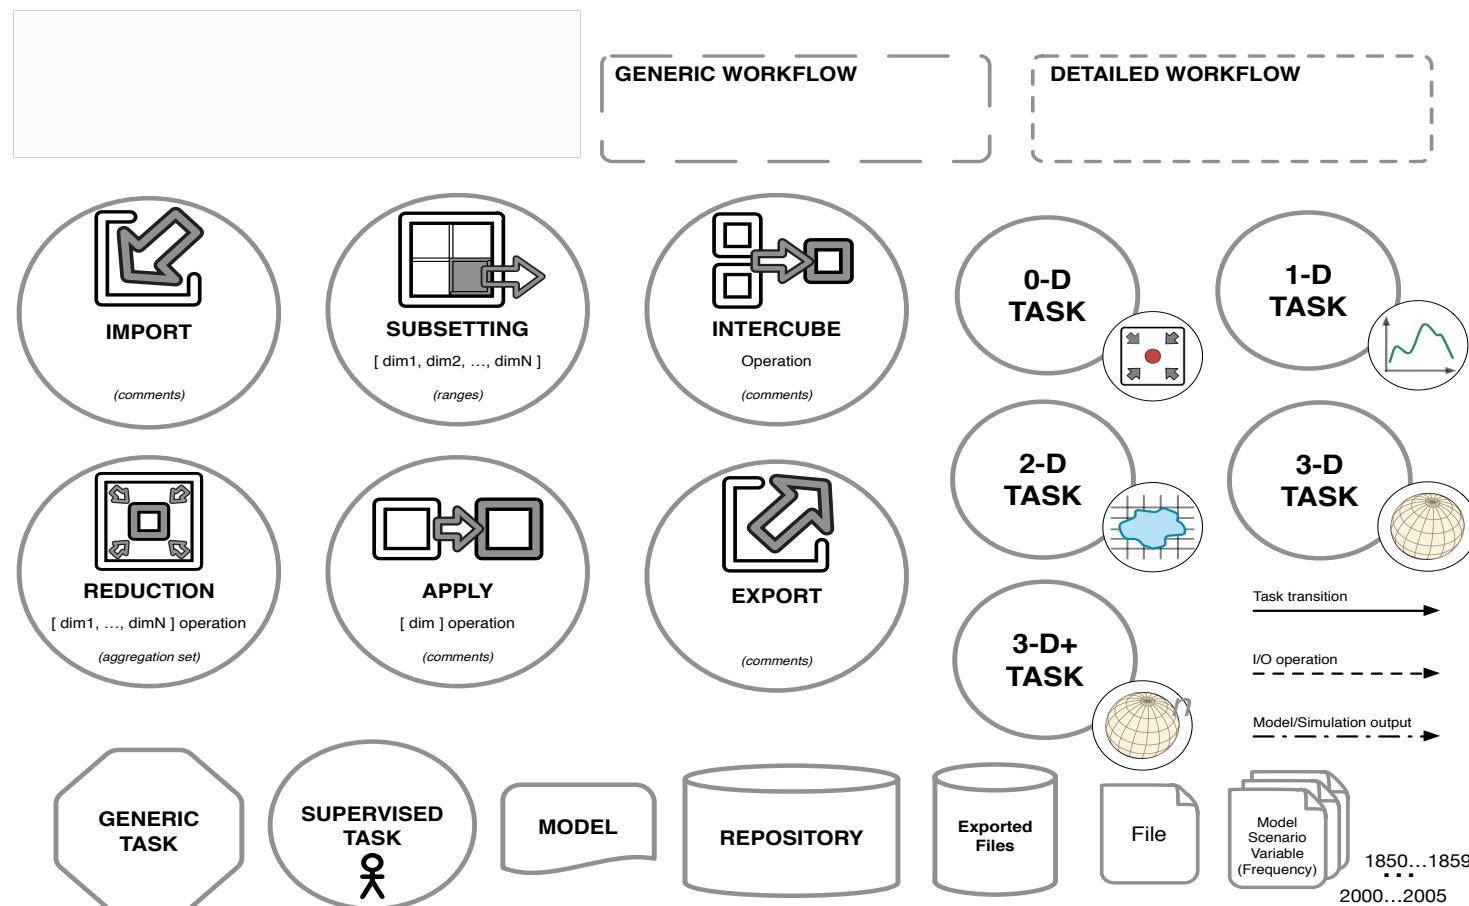
INDIGO - DataCloud

- Time series analysis
- Data subsetting
- Model intercomparison
- Multimodel means
- Massive data reduction
- Data transformation (through array-based primitives)
- Param. Sweep experiments (same task applied on a set of data)
- Climate change signal
- Maps generation
- Ensemble analysis
- Data analytics workflow support
- But also...
- Performance, re-usability, extensibility





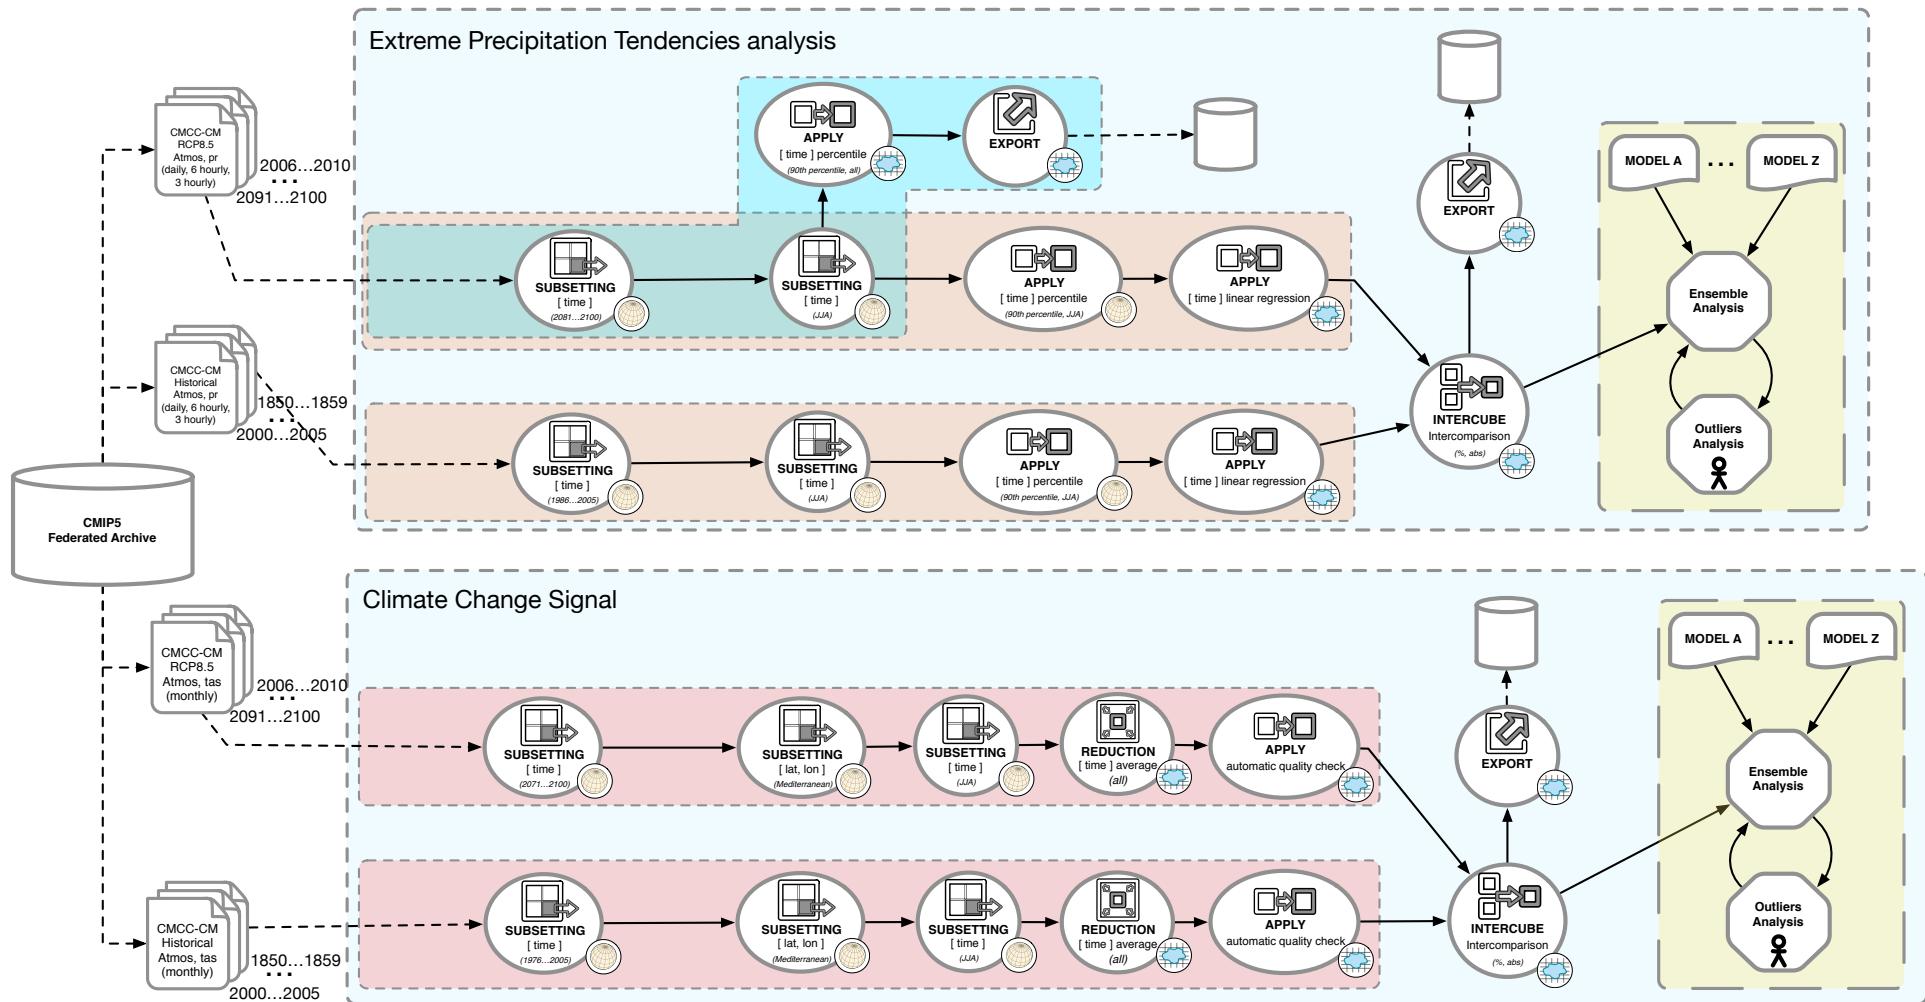
Data Analytics Modeling language





INDIGO - DataCloud

Data analytics use case example

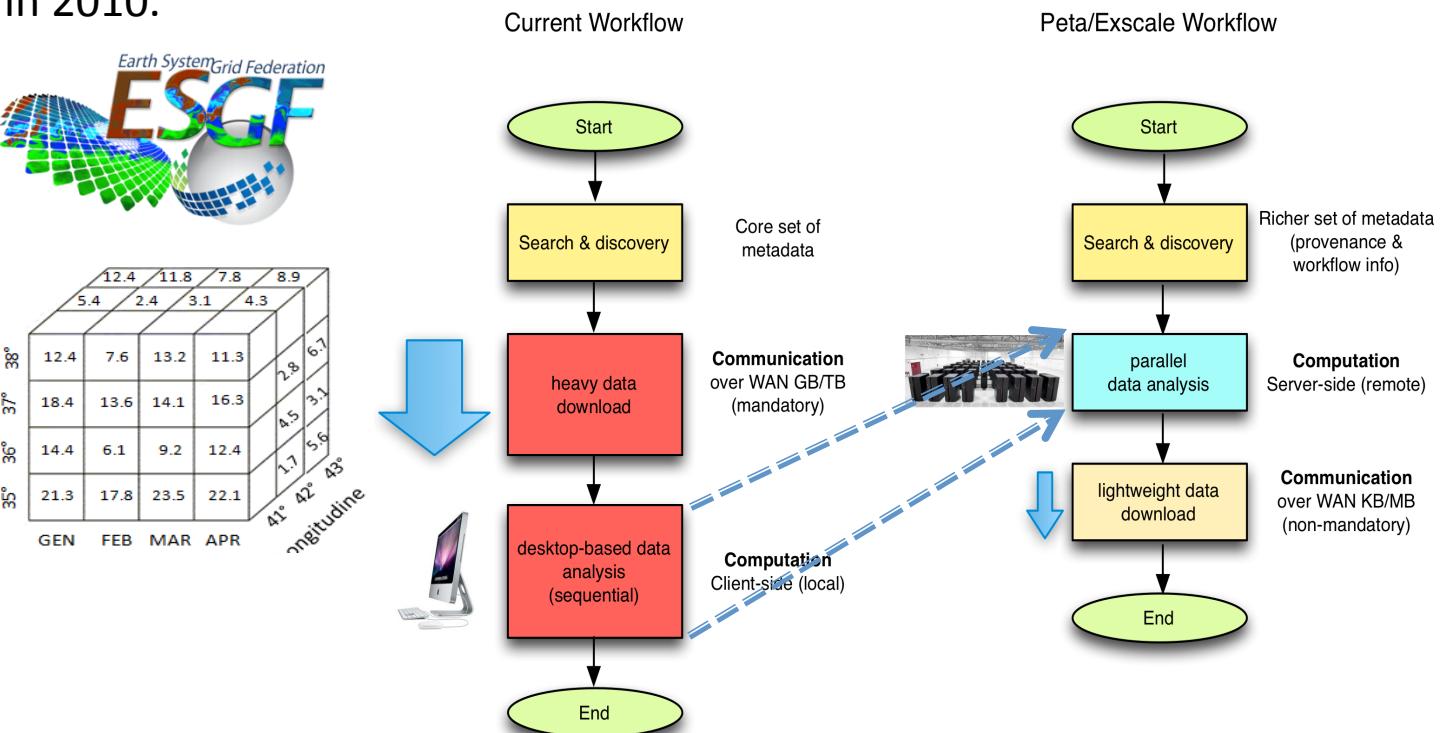
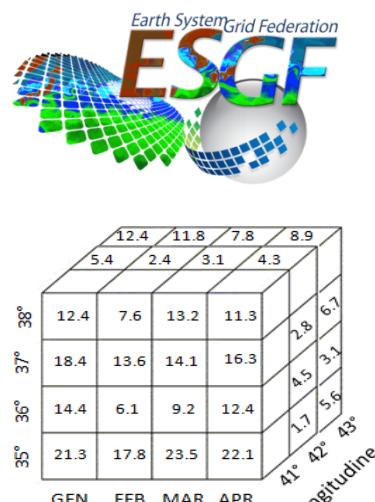
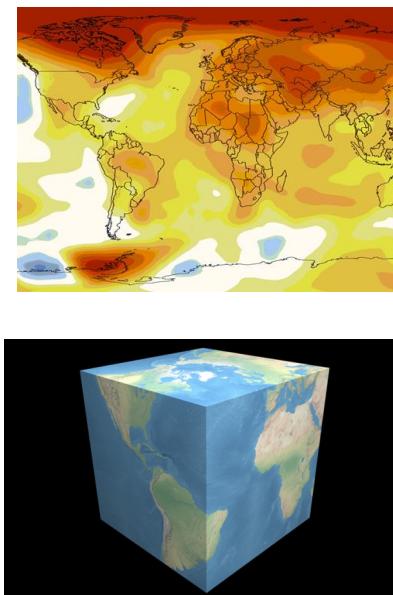




INDIGO - DataCloud

Ophidia

Ophidia is a research effort providing a cross-domain big data analytics framework for the analysis of scientific, multi-dimensional datasets. The framework exploits a server-side, declarative, parallel approach for big data analytics. The design study started in 2009, the implementation stage in 2010.

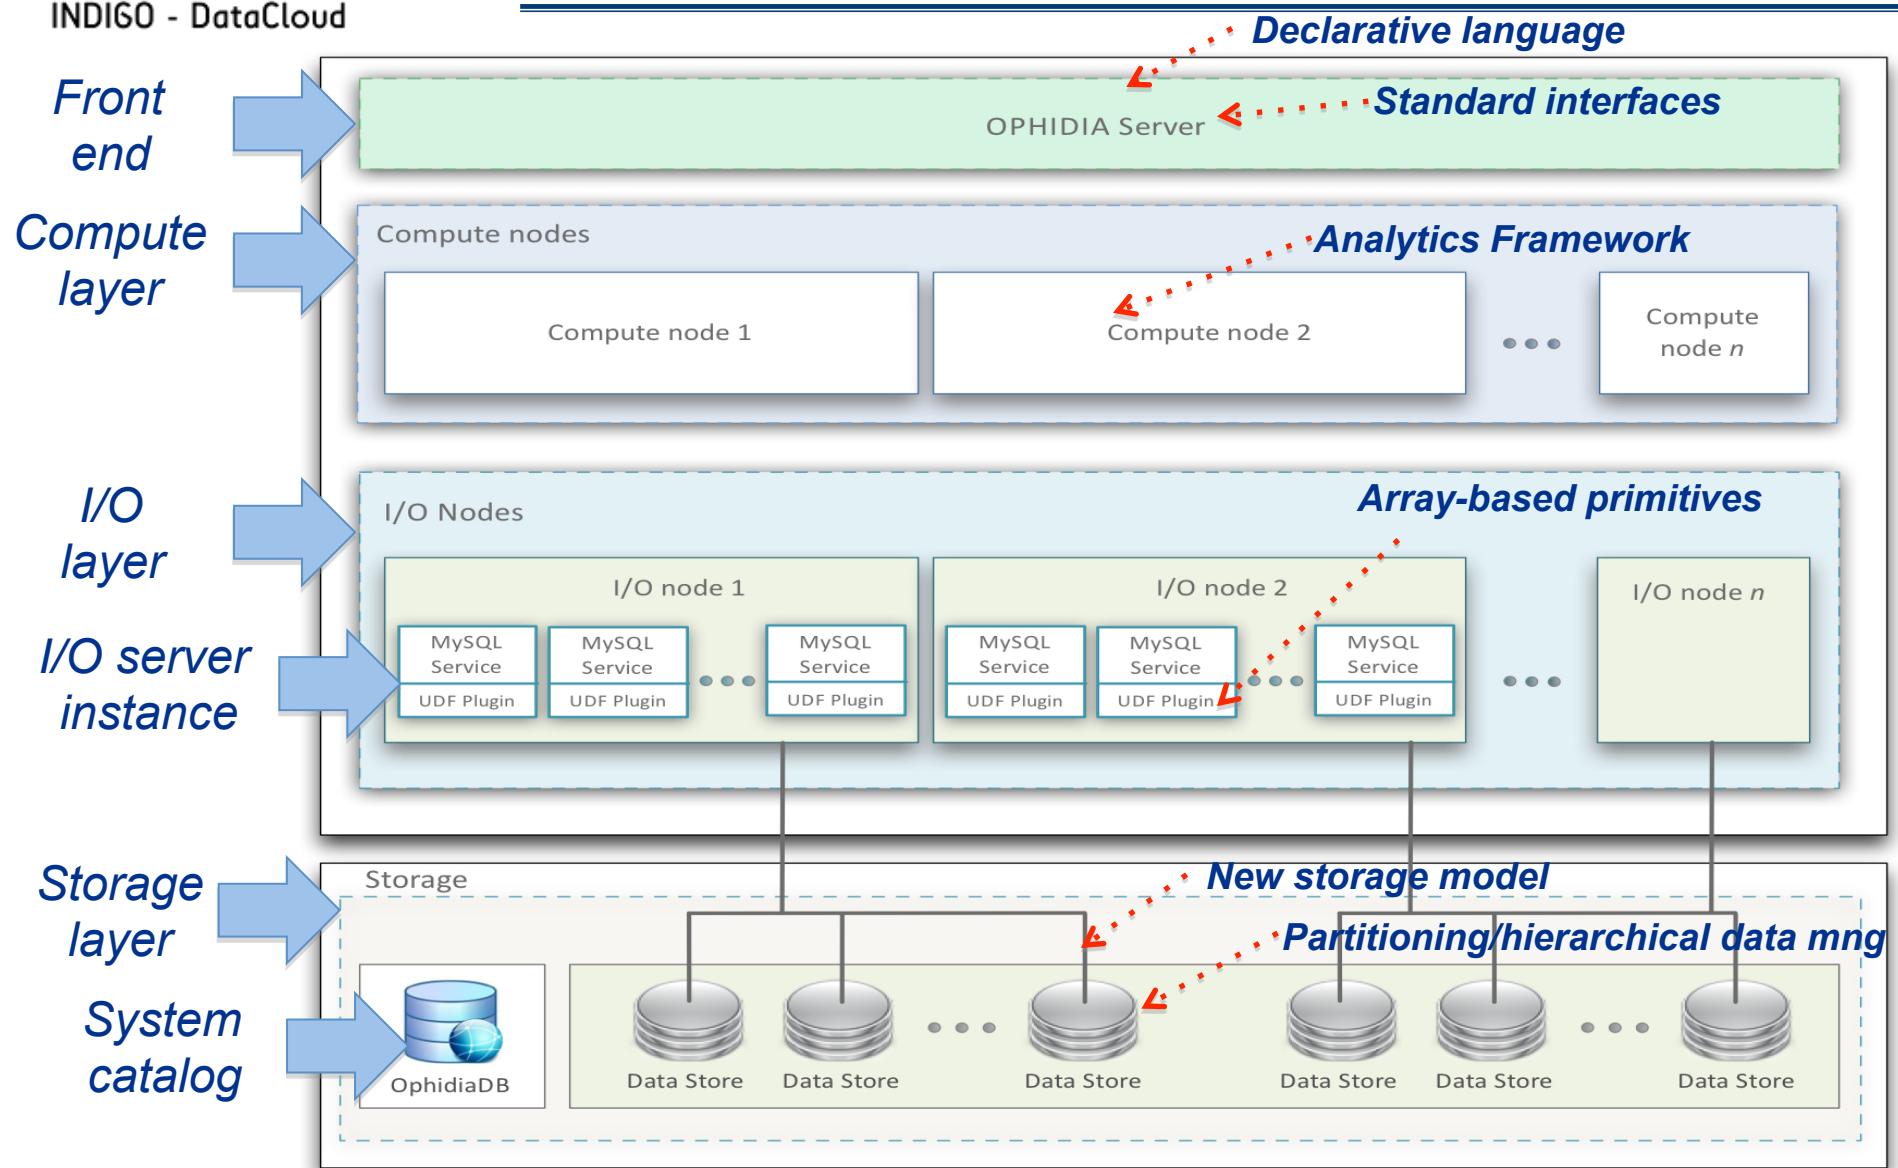


S. Fiore, A. D'Anca, C. Palazzo, I. Foster, D. N. Williams, G. Aloisio, “**Ophidia: toward bigdata analytics for eScience**”, ICCS2013 Conference, Procedia Elsevier, Barcelona, June 5-7, 2013



INDIGO - DataCloud

Ophidia Architecture





INDIGO - DataCloud

The analytics framework: datacube operators (≈ 50)

OPERATOR NAME	OPERATOR DESCRIPTION
Operators “Data processing” – Domain-agnostic	
OPH_APPLY(datacube_in, datacube_out, array based primitive)	Creates the <i>datacube_out</i> by applying the <i>array-based primitive</i> to the <i>datacube_in</i>
OPH_DUPLICATE(datacube_in, datacube_out)	Creates a copy of the <i>datacube_in</i> in the <i>datacube_out</i>
OPH_SUBSET(datacube_in, subset_string, datacube_out)	Creates the <i>datacube_out</i> by doing a sub-setting of the <i>datacube_in</i> by applying the <i>subset string</i>
OPH_MERGE(datacube_in, merge_param, datacube_out)	Creates the <i>datacube_out</i> by merging groups of <i>merge_param</i> fragments from <i>datacube_in</i>
OPH_SPLIT(datacube_in, split_param, datacube_out)	Creates the <i>datacube_out</i> by splitting into groups of <i>split_param</i> fragments each fragment of the <i>datacube_in</i>
OPH_INTERCOMPARISON (datacube_in1, datacube_in2, datacube_out)	Creates the <i>datacube_out</i> which is the element-wise difference between <i>datacube_in1</i> and <i>datacube_in2</i>
OPH_DELETE(datacube_in)	Removes the <i>datacube_in</i>

Data Access
(sequential and parallel operators)

Metadata management
(sequential and parallel operators)

Data processing
(parallel operators, MPI &
OpenMP based)

Import/Export
(parallel operators)

OPERATOR NAME	OPERATOR DESCRIPTION
Operators “Data processing” – Domain-oriented	
OPH_EXPORT_NC (datacube_in, file_out)	Exports the <i>datacube_in</i> data into the <i>file_out</i> NetCDF file.
OPH_IMPORT_NC (file_in, datacube_out)	Imports the data stored into the <i>file_in</i> NetCDF file into the new <i>datacube_out</i> datacube
Operators “Data access”	
OPH_INSPECT_FRAG (datacube_in, fragment_in)	Inspects the data stored in the <i>fragment_in</i> from the <i>datacube_in</i>
OPH_PUBLISH(datacube_in)	Publishes the <i>datacube_in</i> fragments into HTML pages
Operators “Metadata”	
OPH_CUBE_ELEMENTS (datacube_in)	Provides the total number of the elements in the <i>datacube_in</i>
OPH_CUBE_SIZE (datacube_in)	Provides the disk space occupied by the <i>datacube_in</i>
OPH_LIST(void)	Provides the list of available datacubes.
OPH_CUBEIO(datacube_in)	Provides the provenance information related to the <i>datacube_in</i>
OPH_FIND(search_param)	Provides the list of datacubes matching the <i>search_param</i> criteria



Array-based primitives (>100)

Mathematical primitives:

oph_math, oph_mul_scalar, oph_sum_array, oph_sum_scalar, oph_sum_array_r, oph_gsl_sd, oph_gsl_complex_get_abs, oph_gsl_complex_get_arg, oph_gsl_complex_get_imag, oph_gsl_complex_get_real, oph_gsl_complex_to_polar, oph_gsl_complex_to_rect, oph_gsl_dwt, oph_gsl_fft, oph_gsl_idwt, oph_gsl_ifft, oph_gsl_quantile, oph_gsl_sort, oph_gsl_stats, oph_gsl_histogram, oph_gsl_boxplot, oph Petsc_vec_norm, oph Petsc_vec, oph Petsc_vec_r, oph_sum_scalar2, oph_mul_scalar2, oph_compare, ...

Data transformations:

oph_to_bin, oph_to_bit, oph_permute, oph_reverse, oph_dump, oph_convert_d, oph_shift, oph_rotate, oph_concat, oph_predicate, oph_roll_up, oph_get_subarray, oph_get_subarray2, oph_get_subarray3, oph_mask, oph_cast, oph_drill_down (stored procedure), oph_reduce, oph_reduce2, oph_reduce3, oph_aggregate_operator, oph_operator, oph_aggregate_stats, oph_extract, ...

Compression:

oph_compress, oph_uncompress, oph_compress2, oph_uncompress2, oph_id_to_index, oph_size_array, oph_count_array, oph_find, oph_find2, ...

Bit measures processing:

oph_bit_aggregate, oph_bit_import, oph_bit_export, oph_bit_size, oph_bit_count, oph_bit_dump, oph_bit_operator, oph_bit_not, oph_bit_shift, oph_bit_rotate, oph_bit_reverse, oph_bit_subarray, oph_bit_subarray2, oph_bit_concat, oph_bit_find, oph_bit_reduce, ...

Integration of numerical libraries:

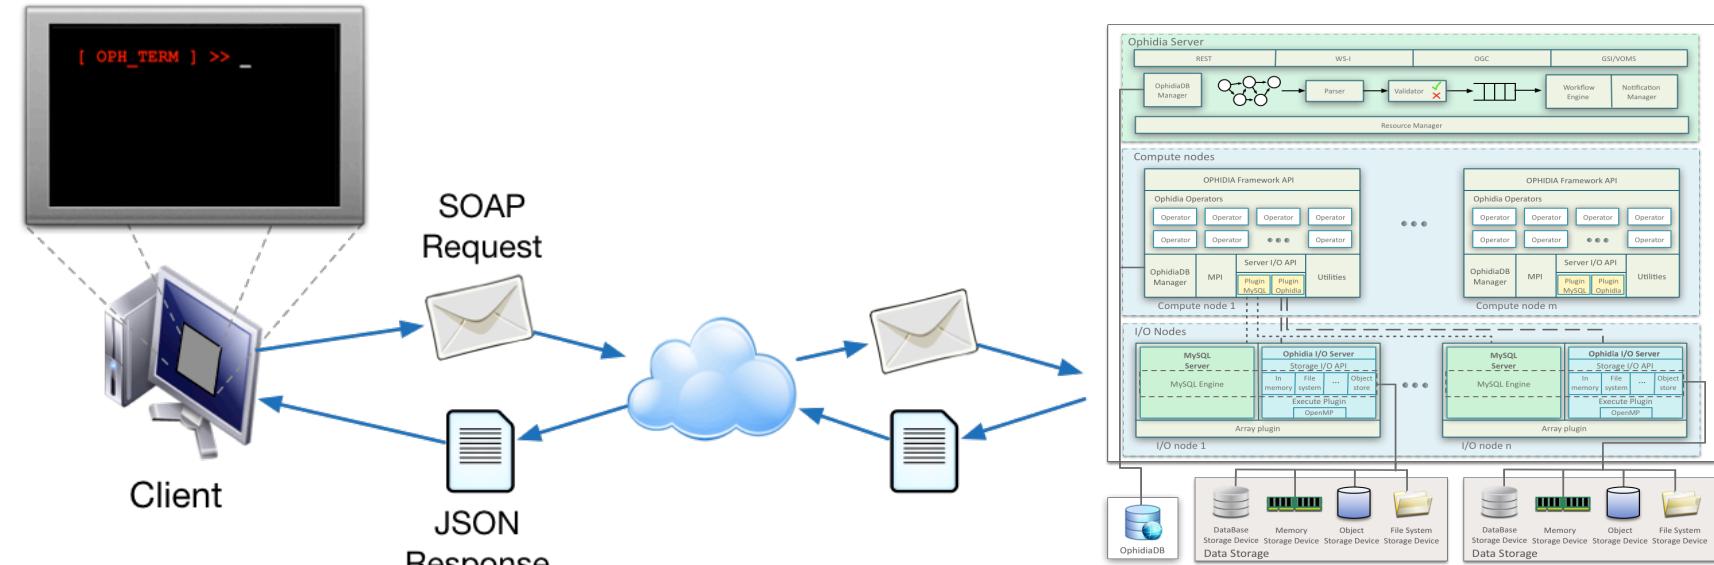
math, GNU GSL, Petsc.



INDIGO - DataCloud

Ophidia terminal

- The Ophidia terminal provides an effective and lightweight way to interact with the Ophidia server
- Bash-like environment (commands interpreter)
- Terminal with history management, auto-completion, specific environment variables and commands with integrated help...
- Easy installation as an only one executable using a small number of well-known and open-source libraries
- More than 15 KLOC
- Simple enough for a novice and at the same time powerful enough for an expert





INDIGO - DataCloud

Ophidia terminal

```
[42..2210] >> oph_cubeschema cube=http://193.204.199.174/demo/18/155;
[Request]:
operator=oph_cubeschema;cube=http://193.204.199.174/demo/18/155;sessionid=http://193.204.199.174/demo/sessions/423749966738743893101423666273632210/experiment;exec_mode=sync;cwd
=/;

[JobID]:
http://193.204.199.174/demo/sessions/423749966738743893101423666273632210/experiment?103#221

[Response]:
Datacube Information
-----
+-----+-----+-----+-----+-----+-----+-----+
| PID | CREATION DATE | MEASURE | MEASURE TYPE | LEVEL | NUMBER OF FRAGMENTS | SOURCE FILE |
+-----+-----+-----+-----+-----+-----+-----+
| http://193.204.199.174/demo/18/155 | 2015-02-11 16:30:02 | T2M | float | 0 | 128 | /home/demo/test_files/test128x128.nc |
+-----+-----+-----+-----+-----+-----+-----+

Datacube Additional Information
-----
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| DESCRIPTION | HOST x CUBE | DBMS x HOST | DATABASES x DBMS | FRAGMENTS x DATABASE | ROWS x FRAGMENT | ELEMENTS x ROW | COMPRESSED | CUBE SIZE | UNIT | NUMBER OF ELEMENTS |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+
| | 1 | 1 | 4 | 32 | 128 | 25 | no | 1.937500 | MB | 409600 |
+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+-----+

Dimension Information
-----
+-----+-----+-----+-----+-----+-----+-----+
| NAME | TYPE | SIZE | HIERARCHY | CONCEPT LEVEL | ARRAY | LEVEL | LATTICE NAME |
+-----+-----+-----+-----+-----+-----+-----+
| lat | double | 128 | oph_base | cell | no | 1 | |
| -----+-----+-----+-----+-----+-----+-----+-----|
| lon | double | 128 | oph_base | cell | no | 2 | |
| -----+-----+-----+-----+-----+-----+-----+-----|
| time | double | 25 | oph_time | 3_hour | yes | 1 | |
+-----+-----+-----+-----+-----+-----+-----+
```



INDIGO - DataCloud

Ophidia terminal

```
[42..2210] >> oph_explorecube cube=http://193.204.199.174/demo/18/192;limit_filter=10
[Request]:
operator=oph_explorecube;cube=http://193.204.199.174/demo/18/192;limit_filter=10;sessionid=http://193.204.199.174/demo/sessions/423749966738743893101423666273632210/experiment;exec_mode=sync;
cwd=/;

[JobID]:
http://193.204.199.174/demo/sessions/423749966738743893101423666273632210/experiment?118#251

[Response]:
T2M
---

+-----+-----+-----+
| lat | lon | T2M |
+-----+-----+-----+
| -90.000000 | -180.000000 | 292.6377563477, 293.7806091309, 292.7806091309, 293.0663452148, 292.9234619141, 293.4949035645, 293.6377563477, 292.9234619141, 292.6377563477, 293.3520507812 |
|-----+-----+-----+
| -90.000000 | -177.187500 | 293.0663452148, 293.6377563477, 293.7806091309, 292.6377563477, 293.6377563477, 293.2091979980, 292.6377563477, 293.2091979980, 293.4949035645, 292.7806091309 |
|-----+-----+-----+
| -90.000000 | -174.375000 | 292.9234619141, 293.2091979980, 293.6377563477, 293.2091979980, 293.3520507812, 293.7806091309, 293.3520507812, 293.2091979980, 293.4949035645, 293.3520507812 |
|-----+-----+-----+
| -90.000000 | -171.562500 | 293.7806091309, 293.0663452148, 293.2091979980, 292.7806091309, 293.4949035645, 293.4949035645, 293.0663452148, 292.9234619141, 293.3520507812, 293.6377563477 |
|-----+-----+-----+
| -90.000000 | -168.750000 | 292.9234619141, 293.4949035645, 293.6377563477, 292.7806091309, 292.6377563477, 293.4949035645, 293.6377563477, 292.7806091309, 293.7806091309, 293.7806091309 |
|-----+-----+-----+
| -90.000000 | -165.937500 | 293.4949035645, 293.4949035645, 292.7806091309, 293.6377563477, 293.4949035645, 293.6377563477, 292.6377563477, 293.0663452148, 293.3520507812, 293.7806091309 |
|-----+-----+-----+
| -90.000000 | -163.125000 | 293.0663452148, 293.6377563477, 293.6377563477, 293.3520507812, 292.6377563477, 293.7806091309, 292.6377563477, 293.6377563477, 292.6377563477, 293.6377563477 |
|-----+-----+-----+
| -90.000000 | -160.312500 | 292.7806091309, 292.7806091309, 293.7806091309, 292.7806091309, 292.7806091309, 292.7806091309, 293.2091979980, 292.6377563477, 292.6377563477, 293.6377563477 |
|-----+-----+-----+
| -90.000000 | -157.500000 | 293.0663452148, 292.6377563477, 293.6377563477, 293.6377563477, 292.7806091309, 292.7806091309, 292.6377563477, 292.9234619141, 292.6377563477, 292.7806091309 |
|-----+-----+-----+
| -90.000000 | -154.687500 | 293.4949035645, 293.7806091309, 293.0663452148, 293.2091979980, 293.0663452148, 293.3520507812, 293.3520507812, 293.0663452148, 293.0663452148, 293.2091979980 |
+-----+-----+-----+
Summary
Selected 10 rows out of 16384
```



INDIGO - DataCloud

Ophidia terminal

```
[42..2210] >> setenv OPH_TERM_VIEWER=dump
[42..2210] >> oph_list
[Request]:
operator=oph_list;sessionid=http://193.204.199.174/demo/sessions/4237499667387438931014236662736322
10/experiment;exec_mode=sync;cube=http://193.204.199.174/demo/18/192;cwd=/;

[JobID]:
http://193.204.199.174/demo/sessions/423749966738743893101423666273632210/experiment?120#255

[Response]:
{
    "response": [
        {
            "objclass": "grid",
            "objkey": "list",
            "objcontent": [
                {
                    "rowvalues": [
                        [
                            [
                                "c",
                                "demo1"
                            ],
                            [
                                "c",
                                "demo2"
                            ]
                        ],
                        "rowfieldtypes": [
                            "string",
                            "string"
                        ],
                        "title": "Ophidia Filesystem: /",
                        "rowkeys": [
                            "T",
                            "PATH"
                        ]
                    }
                }
            ]
        }
    ]
}
```



INDIGO - DataCloud

Ophidia terminal

```
[18..7748] >> oph_list level=3
[Request]:
operator=oph_list;level=3;sessionid=http://localhost/sessions/182736203740511741941409668196547748/experiment;exec_mode=sync;cube=http://localhost/3/11;cwd=/;

[JobID]:
http://localhost/sessions/182736203740511741941409668196547748/experiment?135#267

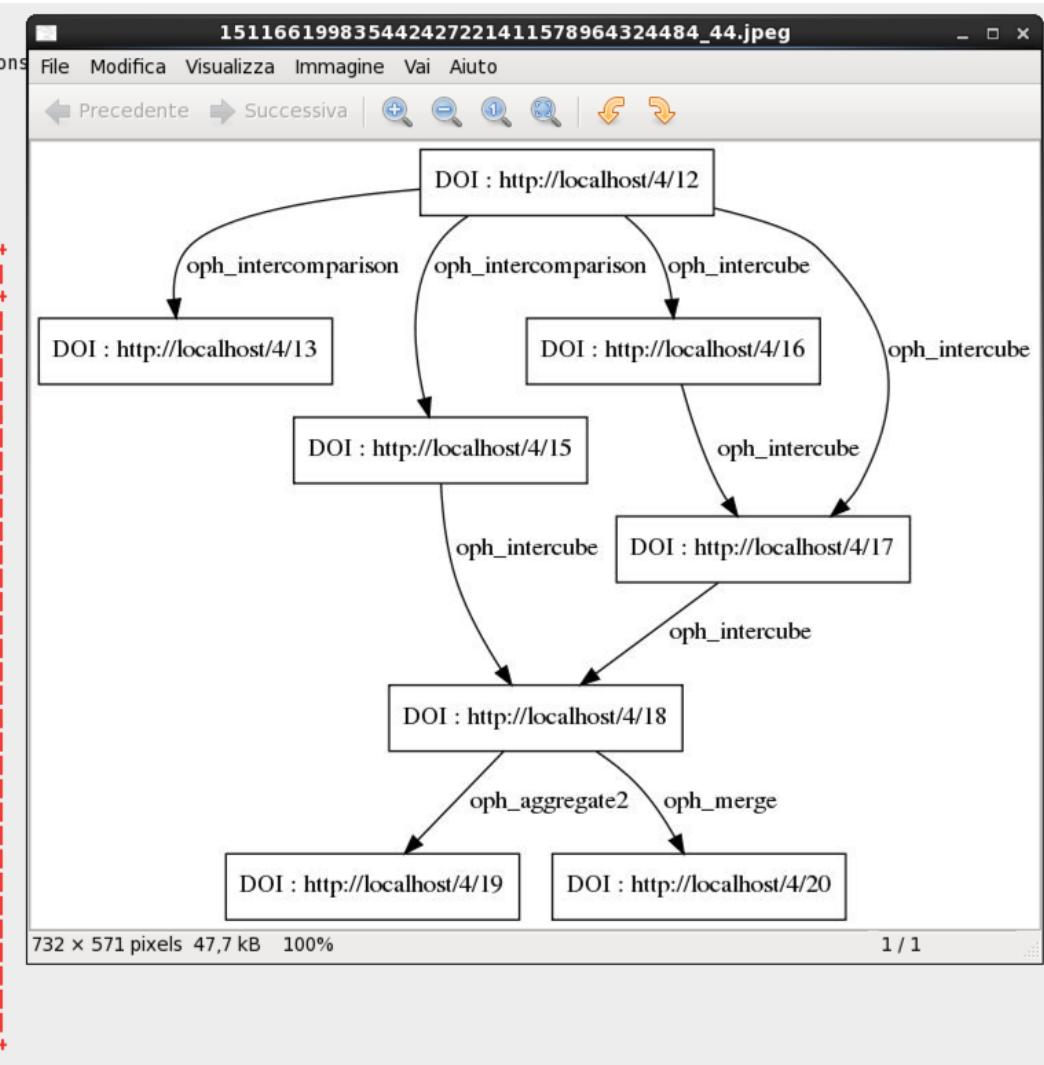
[Response]:
Ophidia Filesystem: /
+---+-----+-----+-----+-----+
| T | PATH      | DATA_CUBE DOI      | MEASURE | LEVEL | SOURCE |
+---+-----+-----+-----+-----+
| f | /myfolder |                         |          |       |          |
| c | /demo     | http://localhost/3/5 | pressure | 0    | RANDOM   |
| c | /demo     | http://localhost/3/6 | pressure | 1    | http://localhost/3/5 |
| c | /demo     | http://localhost/3/7 | pressure | 2    | http://localhost/3/6 |
| c | /demo     | http://localhost/3/8 | pressure | 3    | http://localhost/3/7 |
| c | /demo     | http://localhost/3/9 | pressure | 1    | ...      |
| c | /demo     | http://localhost/3/10 | pressure | 3    | ...      |
| c | /demo     | http://localhost/3/11 | pressure | 2    | ...      |
+---+-----+-----+-----+-----+
```



INDIGO - DataCloud

Ophidia terminal

```
[15..4484] >> oph_cubeio cube=http://localhost/4/12;
[Request]:
operator=oph_cubeio;cube=http://localhost/4/12;sessionid=http://localhost/sessions
[JobID]:
http://localhost/sessions/151166199835442427221411578964324484/experiment?22#43
[Response]:
Cube Provenance
+-----+-----+-----+-----+
| INPUT CUBE | OPERATION | OUTPUT CUBE | SOURCE |
+-----+-----+-----+-----+
|          | ROOT      | http://localhost/4/12 |          |
| http://localhost/4/12 | oph_intercomparison | http://localhost/4/13 |          |
| http://localhost/4/12 | oph_intercomparison | http://localhost/4/13 |          |
| http://localhost/4/12 | oph_intercomparison | http://localhost/4/15 |          |
| http://localhost/4/15 | oph_intercube    | http://localhost/4/18 |          |
| http://localhost/4/18 | oph_aggregate2  | http://localhost/4/19 |          |
| http://localhost/4/18 | oph_merge       | http://localhost/4/20 |          |
| http://localhost/4/12 | oph_intercube    | http://localhost/4/16 |          |
| http://localhost/4/16 | oph_intercube    | http://localhost/4/17 |          |
| http://localhost/4/17 | oph_intercube    | http://localhost/4/18 |          |
| http://localhost/4/18 | oph_aggregate2  | http://localhost/4/19 |          |
| http://localhost/4/18 | oph_merge       | http://localhost/4/20 |          |
| http://localhost/4/12 | oph_intercube    | http://localhost/4/17 |          |
| http://localhost/4/17 | oph_intercube    | http://localhost/4/18 |          |
| http://localhost/4/18 | oph_aggregate2  | http://localhost/4/19 |          |
| http://localhost/4/18 | oph_merge       | http://localhost/4/20 |          |
+-----+-----+-----+-----+
```



Interoperability and standards

- WS-I+ server interface (SSL-based)
- OGC-Web Processing Service (WPS) interface
- RESTful interface (ongoing)
- GSI and VOMS enabled interface (e.g. grid/EGI-FedCloud)
- JSON-schema (core and validation). Draft specification v4
- JSON as interchange format (<http://tools.ietf.org/html/rfc7159>)
- ...

Application-level

- Support for scientific data formats: NetCDF, GRIB, HDF5, GeoTIFF
- Support and domain specific conventions and vocabularies (e.g. CF, CMIP5)

- Contributions to WP2
 - Task T2.1: Research Communities Requirements
 - Task T2.4 Dissemination towards Research Communities
- Contributions to WP6
 - Task T6.1: Libraries and Toolkits
 - Task T6.2 - Science Gateways and User Interfaces
 - Task T6.3: Support for big data driven workflows for e-Science
 - Subtask T6.3.1: Provide scientific workflow support in a “Workflows as a Service” model
 - Subtask T6.3.2: Workflow interface extensions for big data analytics

Task 6.1

- **Task T6.1: Libraries and Toolkits**
 - Development of libraries on top of WP4/WP5 APIs
 - Portlets, servlets, libraries, etc.
 - Focus on scientific data management, workflows, data analytics, data publication, metadata management, data visualization
 - Software layer for Data Analytics Gateways
 - WP2 requirements will be relevant to identify common needs and patterns, prioritize components/libraries, drive the entire process.

■ T6.2.1: Science Gateways

- Science Gateways with a special focus on “*Data Analytics Gateways for eScience*”
 - Driven by the requirements from the WP2 User Communities
- Technologies
 - Portlets, Servlets, Widgets
 - Javascript libraries/frameworks (e.g. ExtJS, Jquery, etc.)
 - Web interactivity (e.g. web-sockets – IETF RFC6455)
 - Visualization libraries (d3.js, amcharts, etc.)
 - ...
- Support for big data analytics functionalities
 - data sub-setting, data reduction/aggregation, statistical data analysis, pivoting, time series, analysis, data summaries concerning multidimensional datasets



INDIGO - DataCloud

Advanced analytics visualization



Task 6.3.1 and Task 6.3.2

- Fine grain “WaaS”, targeting high performance (tightly coupled) data analysis (link to WP5 and related APIs)
- Provide declarative interfaces for big data analytics workflows to support:
 - Massive operators
 - Interactive workflows
 - Interleaved workflows
 - Automated data processing
 - ...and more, based on users requirements.
 - Prioritization will be needed.
- Internals: energy-aware optimizations (equivalence rules, based on analytics operators algebra, at the level of the workflow optimizer)

Demo

([https://www.youtube.com/
user/OphidiaBigData](https://www.youtube.com/user/OphidiaBigData))



INDIGO - DataCloud



www.youtube.com/user/OphidiaBigData



@OphidiaBigData



Data Analytics Terminal:
calling an operator



Data Analytics Terminal:
virtual filesystem navigation



Data Analytics Terminal:
using the manual



Data Analytics Terminal:
sync/async execution mode



Data Analytics Terminal:
simple data manipulation



Data Analytics Terminal:
cube provenance



Data Analytics Terminal:
output rendering



...

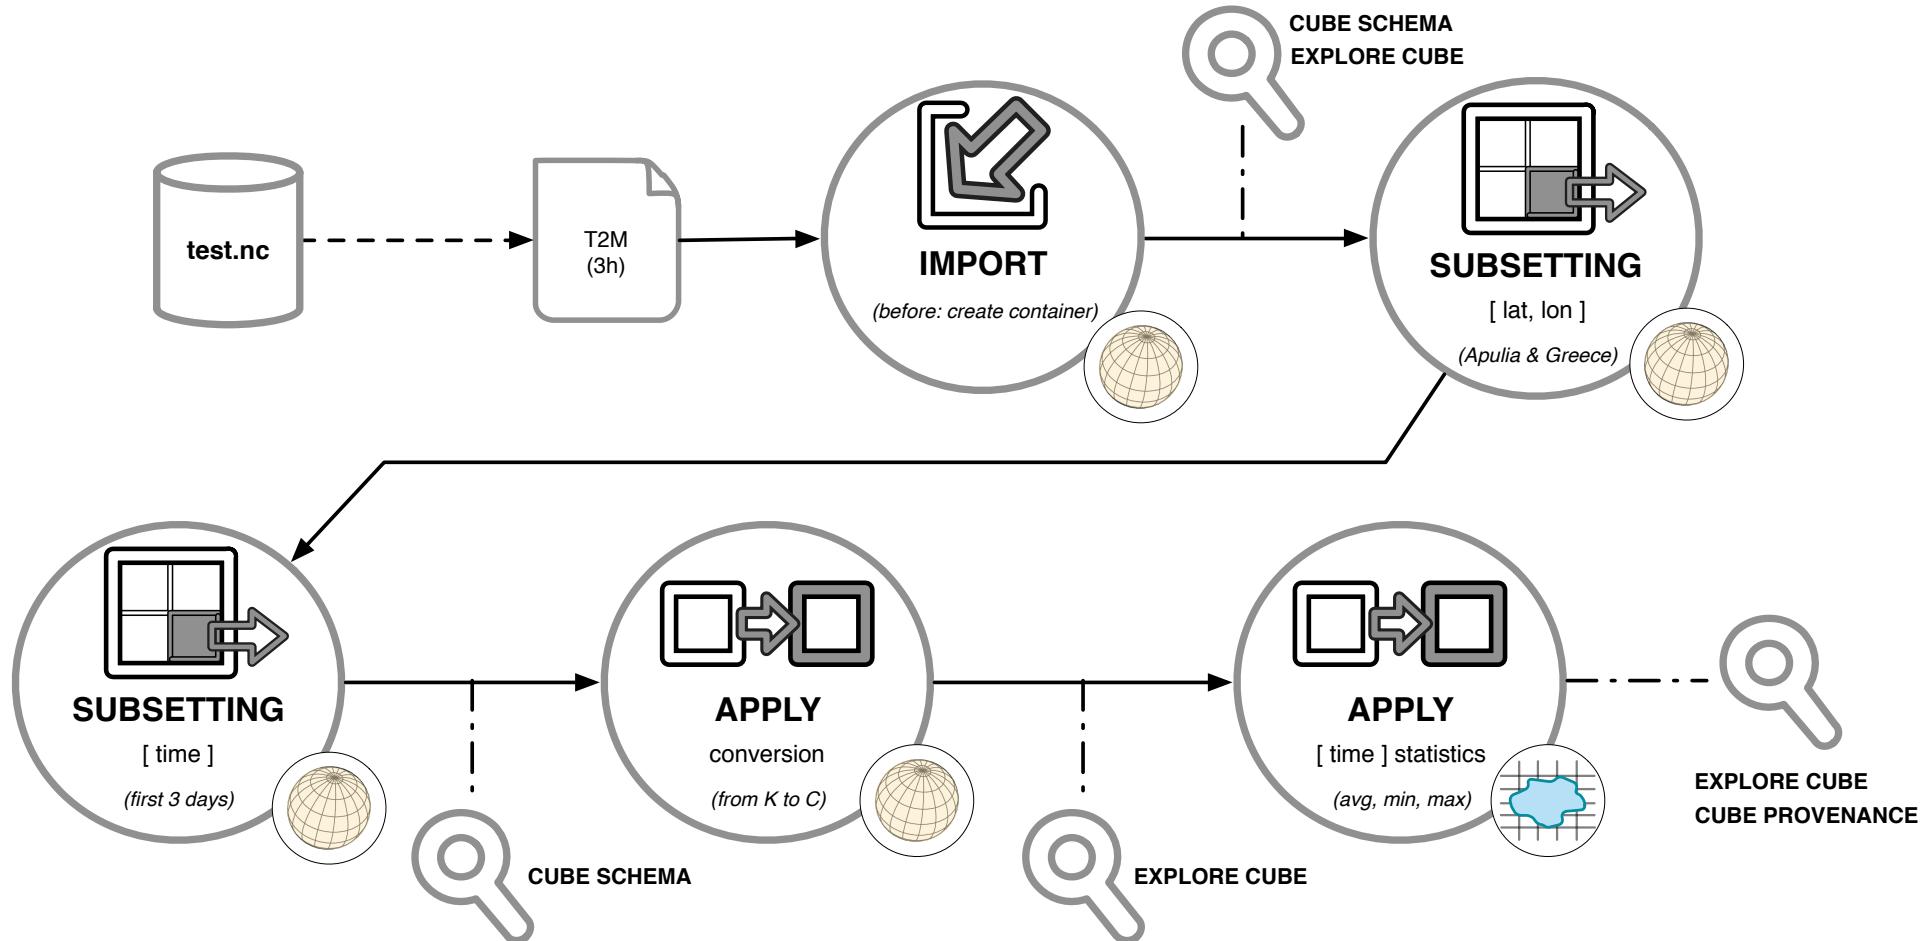


Data Analytics Terminal:
switching between sessions



INDIGO - DataCloud

Temperature conversion & stats



Thank you!
Questions?