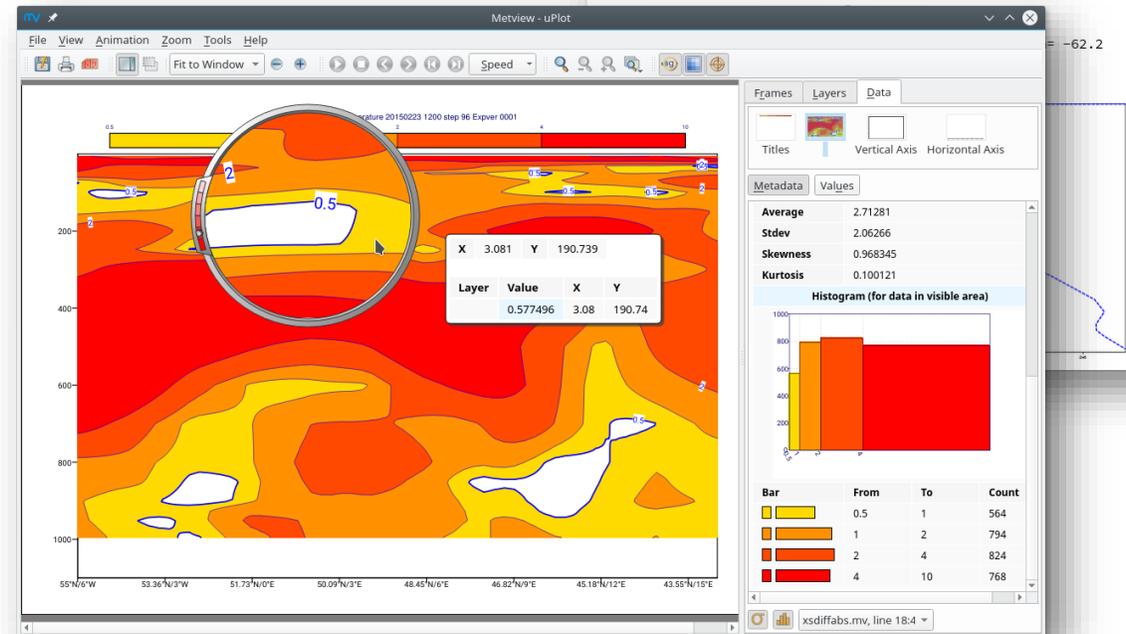


Update on recent developments at ECMWF

Stephan Siemen
Development Section, ECMWF



```
jupyter Vertical profile Last Checkpoint: 3 hours ago (unsaved changes)
File Edit View Insert Cell Kernel Widgets Help
+ < > Run Code
In [11]: import mpy.metview as mpy
        from ipywidgets import interact, interactive, FloatSlider, IntSlider
In [2]: t_fc24 = mpy.read('t_fc24.grib')
In [9]: def plot_vprof_at_point(lat, lon):
        vprofview = mpy.mvertprofview(point = [lat,lon])
        g = mpy.mgraph(graph_line_thickness = 4, graph_line_style = 'dash')
        print('Vertical profile at point: lat=', lat, ' lon=', lon)
        im = mpy.plot(vprofview, t_fc24, g)
        display(im)
        return im
In [10]: interactive_plot = interactive(plot_vprof_at_point, lon=FloatSlider(
        output = interactive_plot.children[-1]
        output.layout.height = '450px'
        output.layout.width = '450px'
        display(interactive_plot)
lat 52.50
```



BOND project – move to Bologna

- Starting in 2019 ECMWF will build its new computing centre in Bologna, Italy
- With this move the whole data centre will be redesigned
- Impact of software developments
 - Review if services can be moved into the cloud (Office 365, GitHub, ...)
 - Migration time is especially challenges – some services will be operated at both sides
 - Removal of old software packages



Evolution of software

To keep a our focus on performance and keep maintenance sustainable we **need to** evolve our development environment and phase out legacy packages

Packages

- SMS → ecFlow
- Ecflowview → ecFlowUI
- Grib_api → ecCodes
- Emoslib/BUFRDC → ecCodes > 2.0
- Emoslib/interpolation → MIR
- BUFR_toolbox → Metview/codes_ui

Languages

- C++ 98 → C++ 14
- Python 2.7 → Python 3.6

Tools

- CMake >2.8 → CMake >3.6

COMING SOON

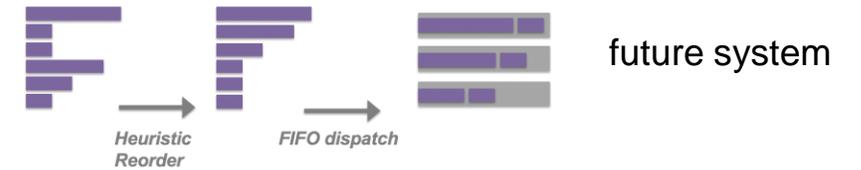
Improving the delivery of forecasts products

- Introduction of new product generation system

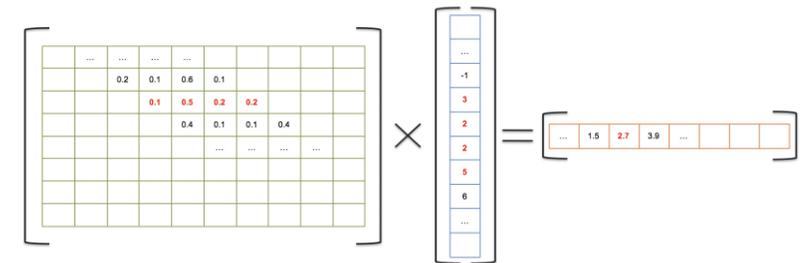
- Already in tests with users
- Move to new system will finish in Q1 2019
- Makes use of new interpolation
- New system will result in easier maintenance, will enable us to grow with new challenges & enable more efficient scheduling of tasks
- Brings MARS archive and dissemination into sync

- New web-based requirements user interface in 2019

- Work progresses well with new interface
- Better real-time validation and feedback on user's changes



New system will allow better scheduling of tasks



Matrix based approach of new interpolation allows for better scaling on new multi-core architectures

```

1 disseminate,
2   stream=enfo,
3   type=pF,
4   number=1 / to / 50,
5   levtype=sfc,
6   param=win,
7   step=0 10fg
8   time=0 10fg3
9   area=- 10fg6
10  grid=0 10fgg5
11         10fgg20
12 disseminat 10fg25
13   stream 10fg1
14   type=c 10app10
15         10app15
16 disseminat 10
17         10na
18         10va
19         10va
20         10va

```

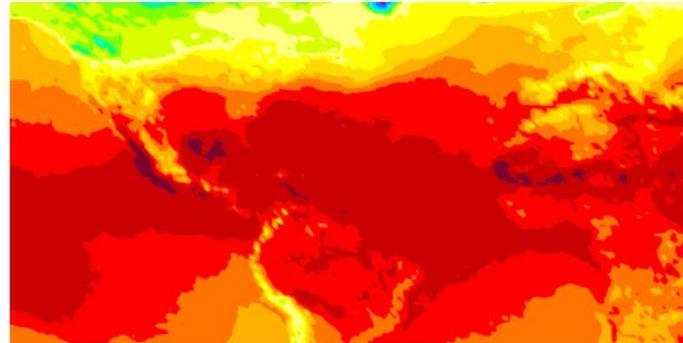
10 metre wind gust since previous post-processing
10 metre wind gust in the last 3 hours
10 metre wind gust in the last 6 hours
10 metre Wind gust of at least 15 m/s
10 metre Wind gust of at least 20 m/s
10 metre Wind gust of at least 25 m/s
10 metre wind gust index
10 metre wind speed
10 metre Wind speed of at least 10 m/s
10 metre Wind speed of at least 15 m/s
10 metre V wind component
10 metre U wind component anomaly
10 metre wind speed anomaly
10 metre V wind component
10 metre V wind component anomaly

The new web interfaces will feature syntax highlighting and auto completion to help user define their requests

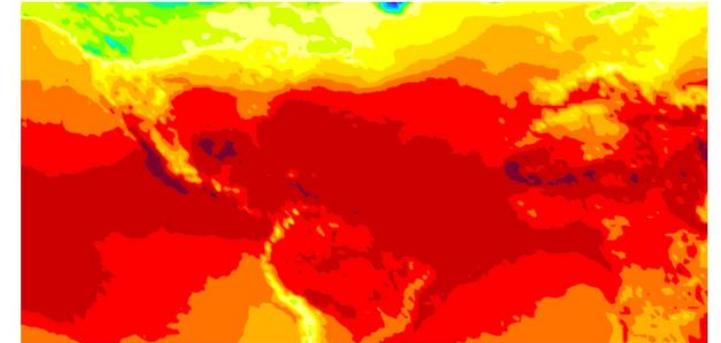
Extensive testing & evaluation has taken place

- New interpolation has different results and therefore needed to be carefully validated
- Much care has been taken to keep technical changes to a minimum for users
- The new system was a chance to review user requirements and adopted them where necessary

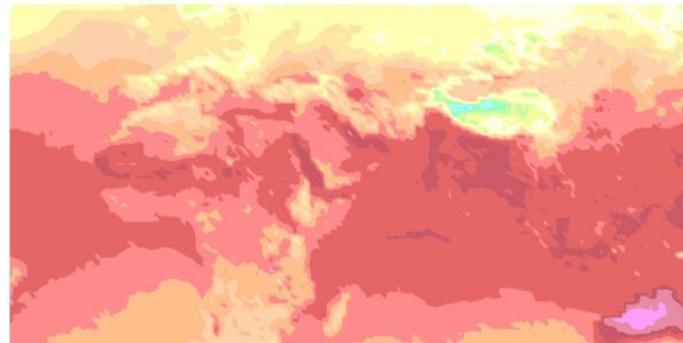
Graphical validation was done through ecCharts, plotting all fields from current (left) and new (right) system.



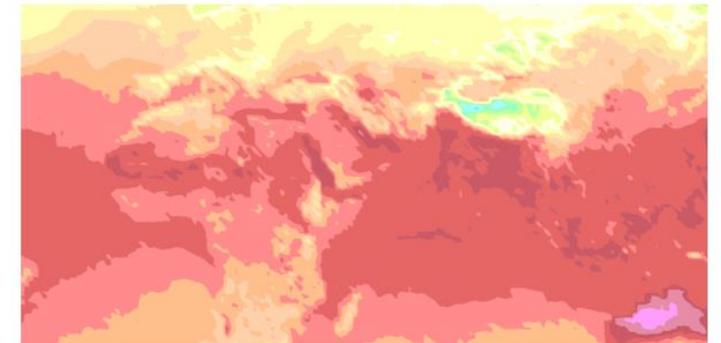
mn2t_sh_all_fm52t48i4: 91883 bytes, 2.08392000198 s



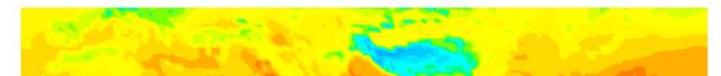
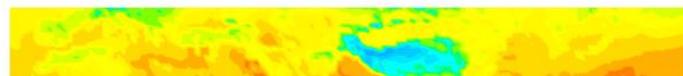
mn2t_sh_all_fm52t48i4: 91883 bytes, 3.53837704659 s, diff: 0



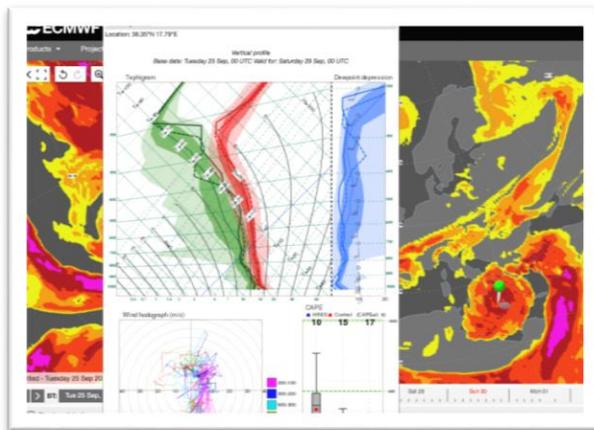
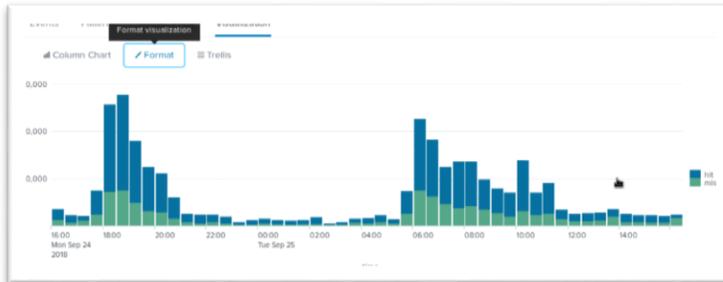
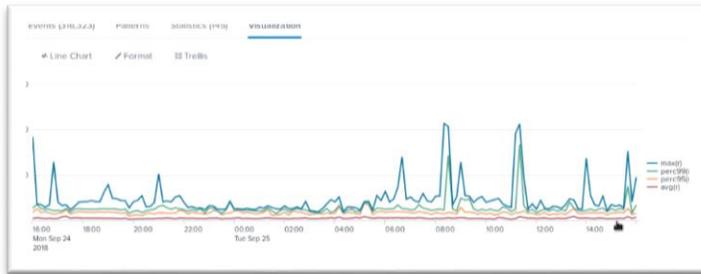
mn2t_sh_all_fm52t48i4_light: 115823 bytes, 3.29667282104 s



mn2t_sh_all_fm52t48i4_light: 115823 bytes, 3.79183411598 s, diff: 0



ecCharts – Improving the user experience



Actual status

- 95% of ecCharts layers are generated in less than 3s.
- Each new request is cached to optimise the response time
- New layers and products are requested by users and added regularly

Plan to improve - ecCharts 2.0

- Improve the user experience by using a 'tile à la google' mechanism (zoom and pan will immediately be more fluid)
- Improve the cache-ability of the requests (the use of tiles will help)
- Improve the 5% of slower requests.
- Improve the deployment procedure to take advantage of cloud technology and new infrastructure.

ecFlowUI is now used in production environment

The image displays five screenshots of the ecFlowUI interface, illustrating its use in a production environment. The screenshots show various views of workflow graphs, node details, and task execution logs.

Top Left Screenshot: Shows a workflow graph with nodes like 'prepare == complete', 'prod', and 'cleanup'. A tooltip for the 'prod' node is visible, showing details such as 'Name: prod', 'Path: /bgen/oper/12/prod', and 'Status: active'. Below the graph is a table of task execution logs.

Node	Status	Type	Trigger	Label	Event	Meter	Status changed
/bgen/oper/12/prod/an/init	complete	task		info: 128.785 pr			2017-Oct-02 1
/bgen/oper/12/prod/an/check	queued	task	postproc == complete	info: errors:			2017-Sep-22 1
/bgen/oper/12/prod/an	active	family					2017-Oct-02 1
/bgen/oper/12/prod/an	active	family					2017-Oct-02 1
/bgen/oper/12/prod	complete	task	prepare == complete	info: OPER12 FC nok			2017-Oct-02 1
/bgen/oper/12/fail	complete	task	prepare eq complete and preparenok	info: 12.OPER.F			2017-Oct-02 1
/bgen/oper/12/cleanup	queued	task	prod == complete	info: 20170820:			2017-Aug-23 1
/bgen/oper/12	active	family		info: 20170820:			2017-Oct-02 1
/bgen/oper/06bc/prod/0090/push	complete	task	/o/main/12/prod2diss==complete and /o/main	info: 19 files, 21			2017-Oct-02 1
/bgen/oper/06bc/prod/0090/postproc	complete	task	postproc == complete	info: 06SCDA.F			2017-Oct-02 1

Top Middle Screenshot: Shows a workflow graph with nodes like 'prepare == complete', 'prod', and 'cleanup'. A tooltip for the 'prod' node is visible, showing details such as 'Name: prod', 'Path: /bgen/oper/12/prod', and 'Status: active'. Below the graph is a table of task execution logs.

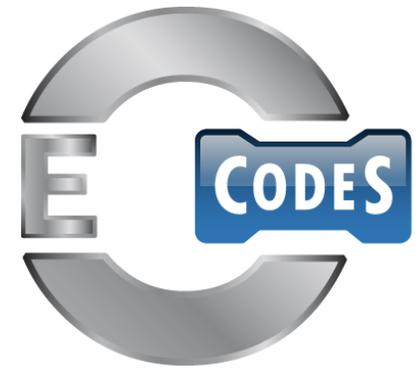
Top Right Screenshot: Shows a workflow graph with nodes like 'prepare == complete', 'prod', and 'cleanup'. A tooltip for the 'prod' node is visible, showing details such as 'Name: prod', 'Path: /bgen/oper/12/prod', and 'Status: active'. Below the graph is a table of task execution logs.

Bottom Left Screenshot: Shows a workflow graph with nodes like 'init', 'postproc', 'check', 'push', 'pgen', and 'ic'. A tooltip for the 'prod' node is visible, showing details such as 'Name: prod', 'Path: /bgen/oper/12/prod', and 'Status: active'. Below the graph is a table of task execution logs.

Bottom Right Screenshot: Shows a workflow graph with nodes like 'make eq complete and /main/YMD le ref/SUITE_STOP', 'limits', 'dummy', 'da_00', 'da_12', 'make', 'setup', 'links', 'rsbias_data_links', 'data_links', 'bins', 'odttools', 'black', 'ifs', 'ssaabs', 'satrad', and 'scat'. A tooltip for the 'bins' node is visible, showing details such as 'Name: bins', 'Path: /bgen/oper/12/prod', and 'Status: active'. Below the graph is a table of task execution logs.

Node	Status	Type	Trigger	Label	Event	Meter	Status changed
/eeda_43r1/main/make/bins/scat	complete	task					2017-Feb-24 18:21:00
/eeda_43r1/main/make/bins/satrad	complete	task					2017-Feb-24 18:27:00
/eeda_43r1/main/make/bins/satmbin	complete	task					2017-Feb-24 18:21:00
/eeda_43r1/main/make/bins/prepdata	complete	task	b2otools == complete				2017-Feb-24 18:53:01
/eeda_43r1/main/make/bins/odttools	complete	task					2017-Feb-24 18:46:00
/eeda_43r1/main/make/bins/obstat	complete	task					2017-Feb-24 18:23:00
/eeda_43r1/main/make/bins/ifs	complete	task	b2otools == complete and black == cor				2017-Feb-24 18:53:01
/eeda_43r1/main/make/bins/black	complete	task					2017-Feb-24 18:22:00
/eeda_43r1/main/make/bins/b2otools	complete	task	odttools==complete and (/e_43r1/main/info: may set 5'				2017-Feb-24 18:51:00

ecCodes – new developments and migrations



- Important to migrate!
 - GRIB_API only supported until end of 2018!
 - BUFREX no further developments
 - ecCodes will be the only package to support new BUFR developments
 - ecCodes can read many BUFR messages BUFREX can't
 - Has many powerful tools and features which BUFREX did not have
- A word on our performance ...
 - We get many reports on performance on ecCodes
 - We treat them as high priority, but it is not always easy to recreate the cases
 - A low level library can be used in different ways – yours might not be the one we use!
 - Please provide us with feedback and examples – and be patient with us
 - We are very thankful for Météo France, MeteoSwiss and DWD and many others for their continuing help
 - As high priority we are looking at improving the threading for GRIB en/decoding
 - After this we will review the BUFR decoding performance

codesUI

- We always had tools to interactively explore GRIB and BUFR data in Metview
- Developed new combined tool for GRIB and BUFR: codesUI
- Allows users to easily access the structure and content of BUFR and GRIB data files
- Available in Metview and standalone
- On ecgate and for download

The screenshot displays the 'temp.bufr - codes.ui' application window. The interface includes a menu bar (File, View, Profiles, Filter, Help), a toolbar with icons for search and refresh, and a 'Key profile: myprof' dropdown. Below this is a navigation area with 'Message: 230', 'Subset: 1', and 'subsets: 1 (number of messages: 420)'. The main area is divided into two panes. The left pane contains a table with columns: Index, E, Type, Sut, C, Sc, Mv, Lv, Ssc, 2, D, T, Lat. The right pane has tabs for 'Data tree', 'Compressed', 'Tables', 'Debug', and 'Locations'. The 'Locations' tab is active, showing a table with columns: Message, Subset, Rank, Latitude, Longitude. Below the table is a world map with numerous red circular markers indicating data locations. At the bottom, a 'Log' panel shows task messages: 'Task: Generating json dump for message: 230', 'Command: /usr/local/apps/eccodes/2.5.0/GNU/5.3.0/bin/bufr_dump -j -w count=230 *var/tmp/cgr/metview/Local/BUFR/ECC/temp.bufr', 'Task: Parsing JSON dump for message: 230', and 'Message: 230 Generating json dump : DONE'.

Index	E	Type	Sut	C	Sc	Mv	Lv	Ssc	2	D	T	Lat
211	3	2	101	ecmf	0	13	1	1	1	20081208	120000	45.42
212	3	2	101	ecmf	0	13	1	1	1	20081208	120000	43.07
213	3	2	101	ecmf	0	13	1	1	1	20081208	120000	43.33
214	3	2	101	ecmf	0	13	1	1	1	20081208	120000	40.7
215	3	2	101	ecmf	0	13	1	1	1	20081208	120000	39.72
216	3	2	101	ecmf	0	13	1	1	1	20081208	120000	37.4
217	3	2	101	ecmf	0	13	1	1	1	20081208	120000	36.05
218	3	2	101	ecmf	0	13	1	1	1	20081208	120000	32.63
219	3	2	101	ecmf	0	13	1	1	1	20081208	120000	22.32
220	3	2	101	ecmf	0	13	1	1	1	20081208	120000	38.25
221	3	2	101	ecmf	0	13	1	1	1	20081208	120000	37.97
222	3	2	101	ecmf	0	13	1	1	1	20081208	120000	36.03
223	3	2	101	ecmf	0	13	1	1	1	20081208	120000	35.12
224	3	2	101	ecmf	0	13	1	1	1	20081208	120000	34.68
225	3	2	101	ecmf	0	13	1	1	1	20081208	120000	33.28
226	3	2	101	ecmf	0	13	1	1	1	20081208	120000	37.1
227	3	2	101	ecmf	0	13	1	1	1	20081208	120000	-0.68
228	3	2	101	ecmf	0	13	1	1	1	20081208	120000	54.75
229	3	2	101	ecmf	0	13	1	1	1	20081208	120000	51.12
230	3	2	101	ecmf	0	13	1	1	1	20081208	120000	52.4
231	3	2	101	ecmf	0	13	1	1	1	20081208	120000	50.4
232	3	2	101	ecmf	0	13	1	1	1	20081208	120000	48.03
233	3	2	101	ecmf	0	13	1	1	1	20081208	120000	49.8
234	3	2	101	ecmf	0	13	1	1	1	20081208	120000	47.92
235	3	2	101	ecmf	0	13	1	1	1	20081208	120000	-67.57
236	3	2	101	ecmf	0	13	1	1	1	20081208	120000	-3.85
237	3	2	101	ecmf	0	13	1	1	1	20081208	120000	-6.77
238	3	2	101	ecmf	0	13	1	1	1	20081208	120000	-9.38
239	3	2	101	ecmf	0	13	1	1	1	20081208	120000	-13.27
240	3	2	101	ecmf	0	13	1	1	1	20081208	120000	-17.73
241	3	2	101	ecmf	0	13	1	1	1	20081208	120000	-37.67
242	3	2	101	ecmf	0	13	1	1	1	20081208	120000	-20.27
243	3	2	101	ecmf	0	13	1	1	1	20081208	120000	-23.5
244	3	2	101	ecmf	0	13	1	1	1	20081208	120000	2.48
245	3	2	101	ecmf	0	13	1	1	1	20081208	120000	-2.43
246	3	2	101	ecmf	0	13	1	1	1	20081208	120000	-3.15

Message	Subset	Rank	Latitude	Longitude
226	1	-1	37.1000	127.0300
227	1	-1	-0.6800	73.1500
228	1	-1	54.7500	17.5300
229	1	-1	51.1200	16.8800
230	1	-1	52.4000	20.9700
231	1	-1	50.4000	30.5700
232	1	-1	48.0300	33.2200
233	1	-1	49.8000	92.0800
234	1	-1	47.9200	106.8700

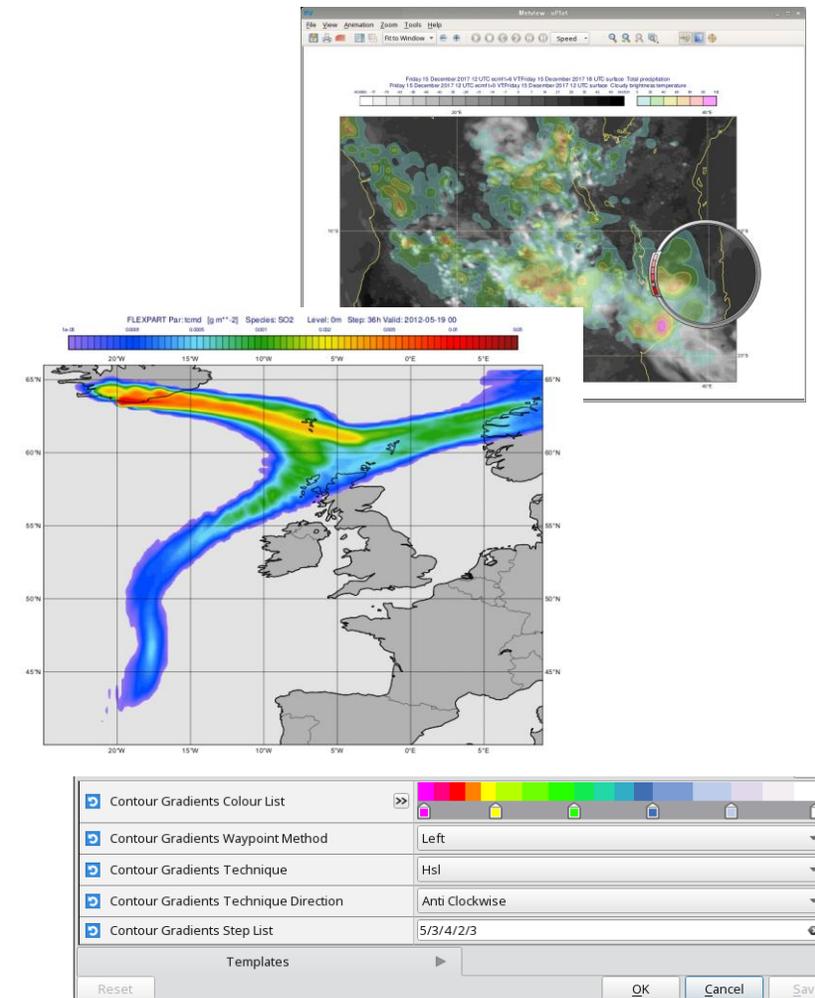
Magics 3 & Metview 5

- **Magics 3.0**

- Mature version of library
- Biggest change – new Python interface for version 2 & 3
- Better automatic support for NetCDF

- **Metview 5.0**

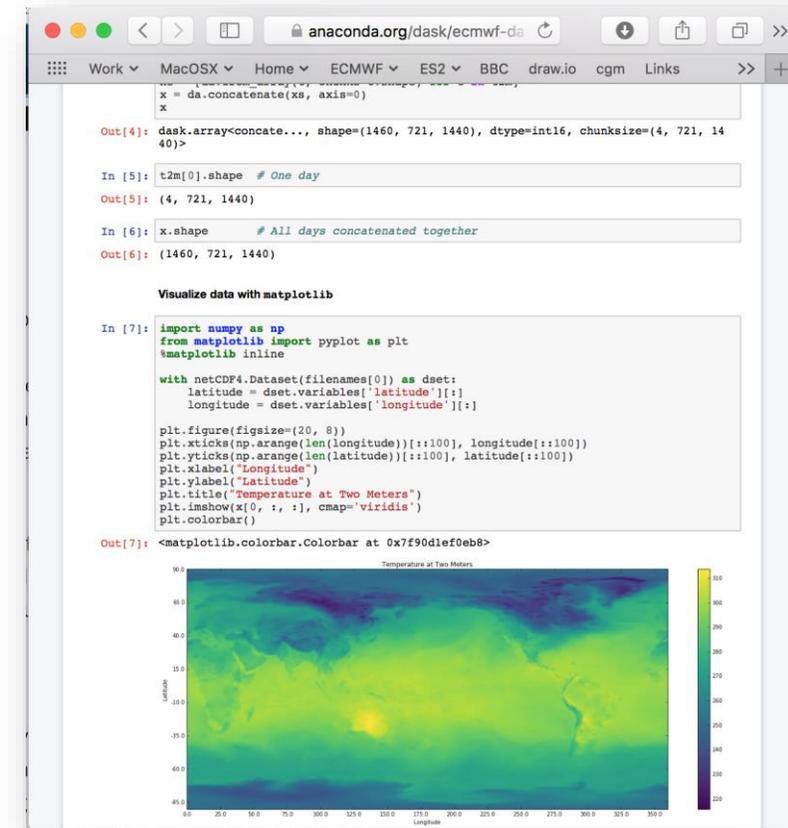
- Motif interface is removed & support Qt5
- New features for interactive editing of plots & new colour schemes
- New interface to Flexpart
- First version to support new Python interface ...
- <https://software.ecmwf.int/wiki/display/METV/Version+5.0+Updates>



→ We will increase major version number more frequently than in the past

Building a Python framework to work with ECMWF data

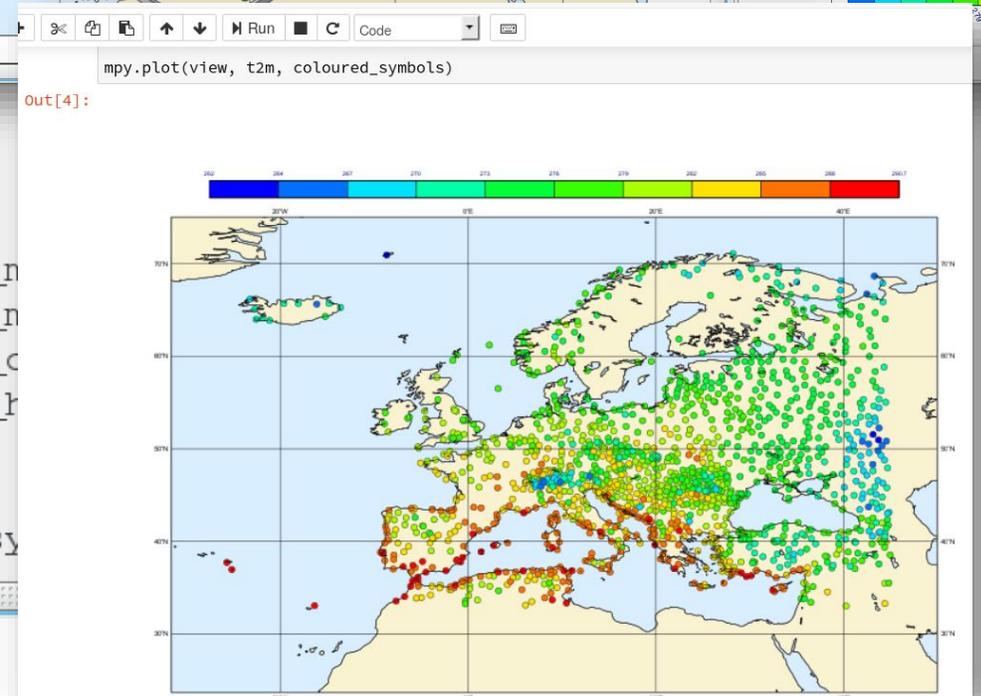
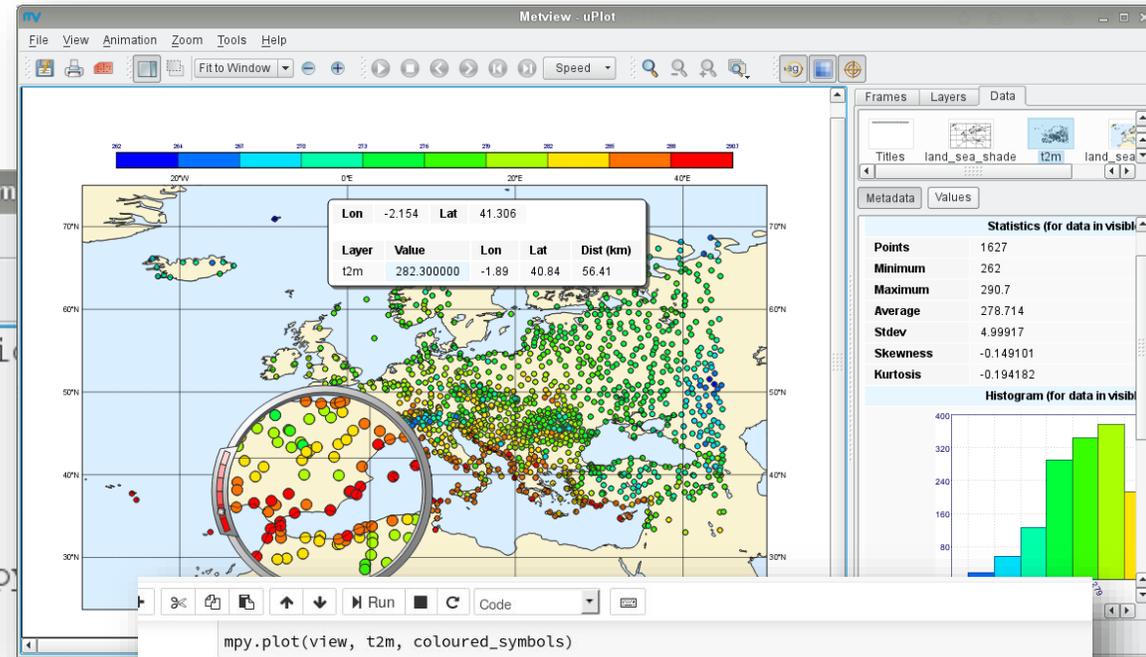
- Important target to fully support Python 3 by the end of 2018
 - Review how we use Python – best practices on how to write scripts
 - Make use of all the possibilities the Python3 eco-system offers
 - Offer JupyterLab based environment for users
- We are making progress
 - Metview-Python already based on Python 3.6
 - Magics & webapi offer Python 2 and 3 interface
 - We created a Python 3 interface to load GRIB data into xarray
 - Using ecCodes; but high-level data structure for NetCDF-CF
 - Used by Metview-Python & CDS toolbox
- We want to progress from the first generation Python interfaces we offered
 - Much better integrated with PyData eco-system and easy installable through pip & conda



The Metview Python framework

- Allow users of Metview to use easily the power of Python but still have all functionality of Metview; including visualisation
- Users can continue in their familiar environments or work in the Python Jupyter environments

```
Python Script.py* - /home/graphics/cgi/m
File Edit View Insert Program Settings
18 map_area_definitio
19 area
20 coastlines
21 )
22
23 coloured_symbols = mpy
24 legend
25 symbol_type
26 symbol_table_mode
27 symbol_outline
28 symbol_outline_colour
29 symbol_advanced_table_n
30 symbol_advanced_table_n
31 symbol_advanced_table_c
32 symbol_advanced_table_h
33 )
34
35 plot(view, t2m, coloured_sy
```



Importance to engage with community

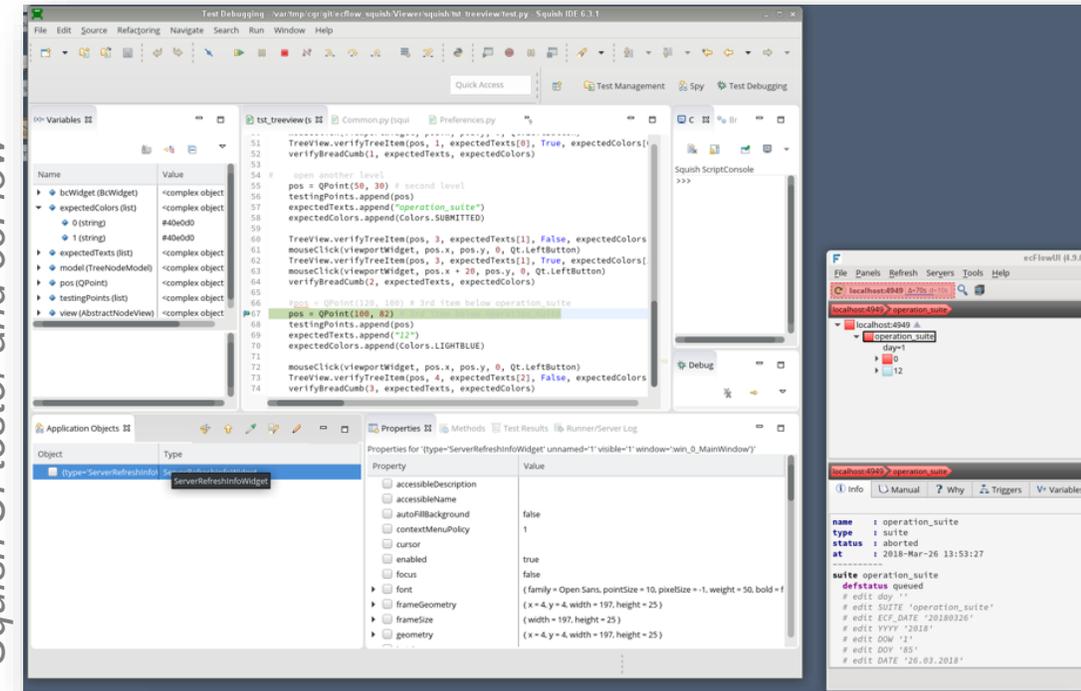
- There are already many good efforts and solutions out there
 - Python seems to be the language of choice
- ECMWF hosted a two day workshop for developers of Python frameworks for earth sciences
 - Many good “wheels” which do not need to be reinvented
 - Building a community is more than just releasing software under Open Source
 - ‘Open Source’ versus ‘Open Development’
- Continue outreach
 - 18-19 October 2nd Python frameworks workshop



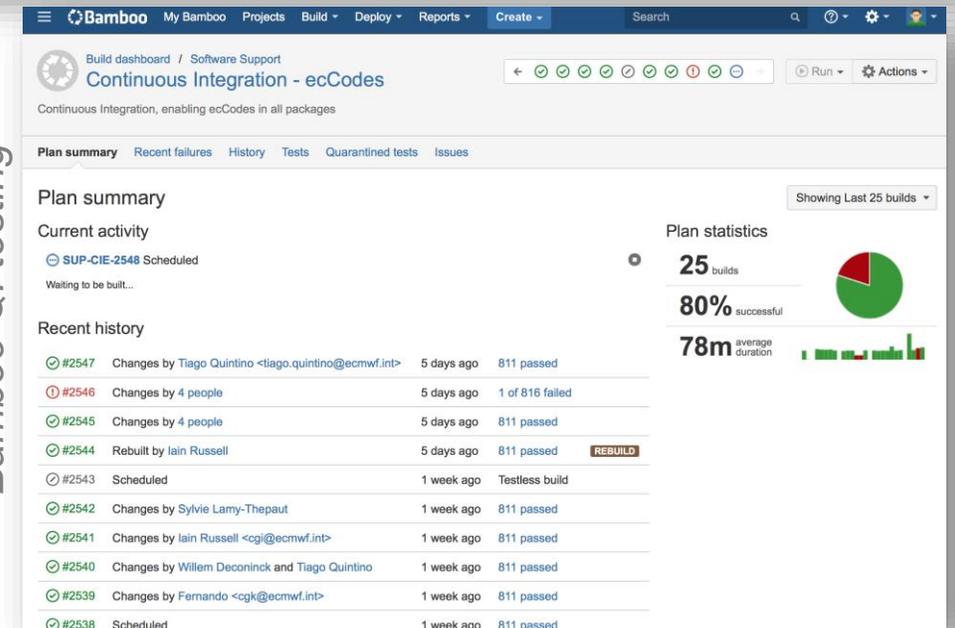
Software development infrastructure

- ECMWF had a series of consultancies to improve QI/UI testing, software installations and make better use of new language features
- Want to engage more with communities which can provide binary packages of ECMWF software
 - Possibly make public git repos available
- ECMWF embraces *Open Development*
 - Started ECMWF space on GitHub
 - Looking at contributing to third-party open source projects, rather than develop own packages

Squish UI tester and ecFlow



Bamboo QI testing



Learning in focus

eLearning

We have completed the following modules:

- *Batch system on ECGATE*
- *ecCodes (GRIB, BUFR)*
- *ecFlow*
- *MARS*
- *Metview*

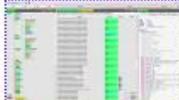
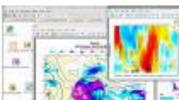
Two more modules in the pipelines:

- *Compiling on ECGATE*
- *ecCodes: advanced GRIB decoding*

Follow us on #ECLearn

Software and Computing services

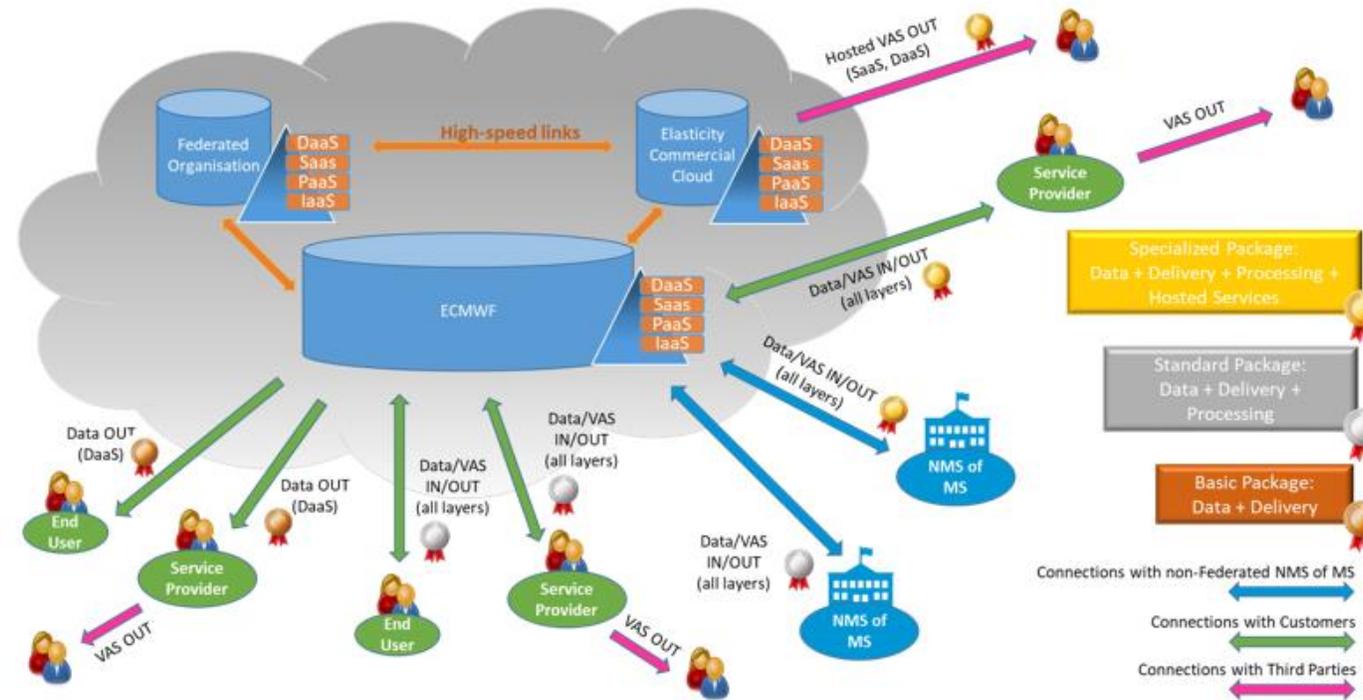
A collections of online resources dedicated to ECWMF software, applications and computing services

 ecCodes: decoding with GRIB tools 1 hour	 ecCodes: advance features (TBA) 1 hour	 Introduction to BUFR decoding with ecCodes 1 hour
 MARS - ECMWF's meteorological archive 1 hour	 MARS - advanced retrievals, data manipulation and computations 1 hour	 A starter guide to ecFlow 1 hour
 Introduction to Metview 1 hour	 Metview for the single-column model (SCM) 20 minutes	 Using ECMWF computing facilities: the batch system 1 hour

<https://www.ecmwf.int/en/learning/education-material/elearning-online-resources>

ECMWF – web services & the cloud

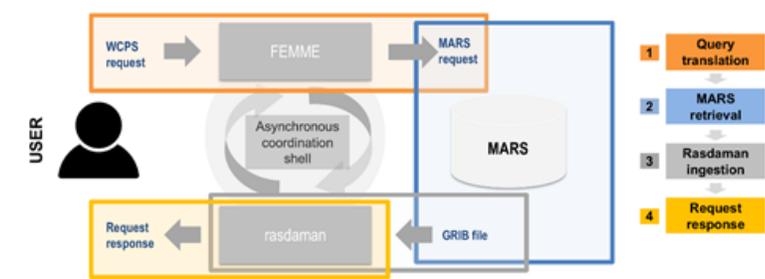
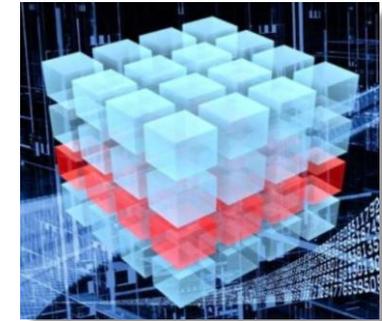
- ECMWF looks together with its partners on providing private clouds
 - European Weather Cloud with EUMETSAT & DIAS WEKEO
- Looking at better opportunities for users to make use of ECMWF forecast data



H2020 EarthServer-2 project



- Came to an end in April after three years
- Explored DataCube technology for fast retrievals of time series
 - We hoped to provide interface for users to build their own applications
- We learned a lot
 - on OGC standards and services and how they could help our users
 - on how we could provide web services to enable uses to build their applications
 - on how to work with commercial third-party software providers
- We built a demonstration WCS/WCPS service
 - Tried to ingest data directly and connect to MARS archive
 - Trialled software did not meet requirements and service is now switched off
 - Lessons learned are integrated in new developments



EarthServer-2 has received funding from the European Union's Horizon 2020 Research and Innovation programme under Grant Agreement no. 671951

ECMWF Summer of Weather Code (ESoWC)



Software developers, start-ups, weather professionals and scientists, working with ECMWF mentors to create new and innovative weather-related code.

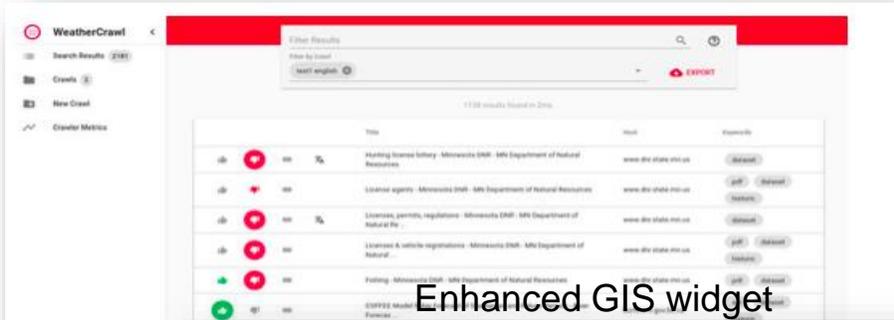
The new developments will help ECMWF support national weather services in its Member and Co-operating States.

5 teams – 5 challenges

- Enhanced GIS widget
- Web crawler for hydrological data
- Migration of calibration software to Python
- Innovative visualisations
- NETCDF as a file system

ECMWF Summer of Weather Code (ESoWC)

WEB crawler

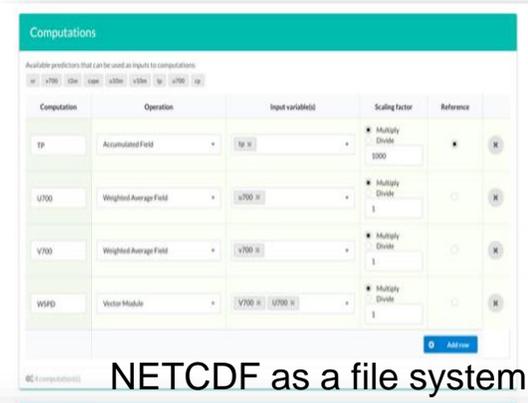


Enhanced GIS widget

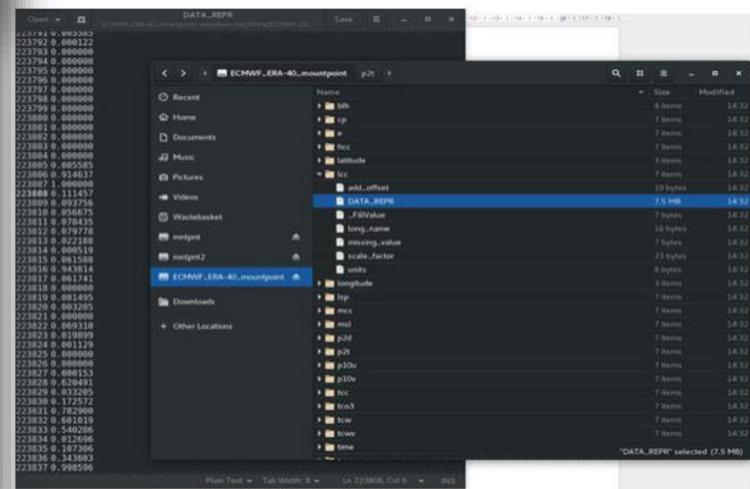


Enhanced GIS widget to select and display areas on a map.

Migration calibration software

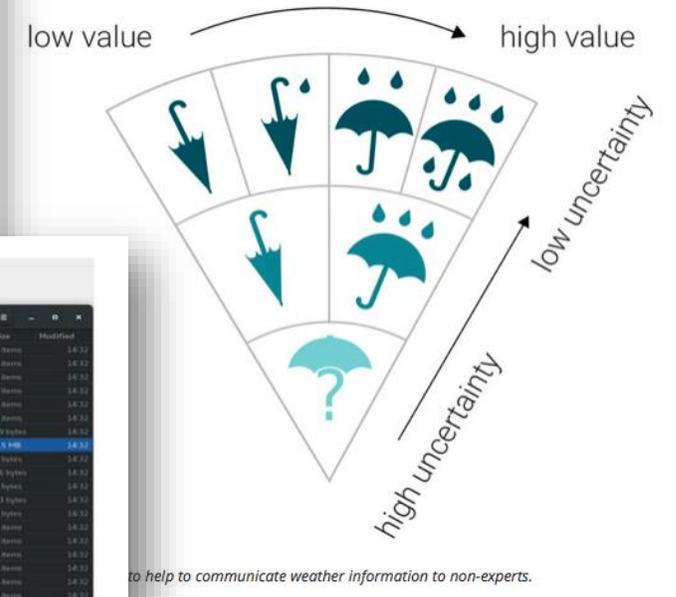


NETCDF as a file system



New software which allows users to easilmount, view, and edit the contents of a NetCDF dataset.

Innovative visualizations



to help to communicate weather information to non-experts.