

MARS – Introduction and basic concepts

MS Computing training course

Bojan Kasic,

Carsten Maass, and Dominique Lucas

Contents

- **Introduction**
- **MARS architecture**
- **Meteorological content**
- **MARS language**
- **Retrieving data**
- **Practical session**


Meteorological Archival and *Retrieval* System

- Meteorological data (GRIB: fields, BUFR, ODB: observations)
- Large amount of data (size of archive & number of fields)
- Operational & Research environment
- Batch & interactive modes
- Large number of users with different requirements:
 - large datasets rarely ↔ few fields very often
- Heterogeneous environment



Introduction – MARS components

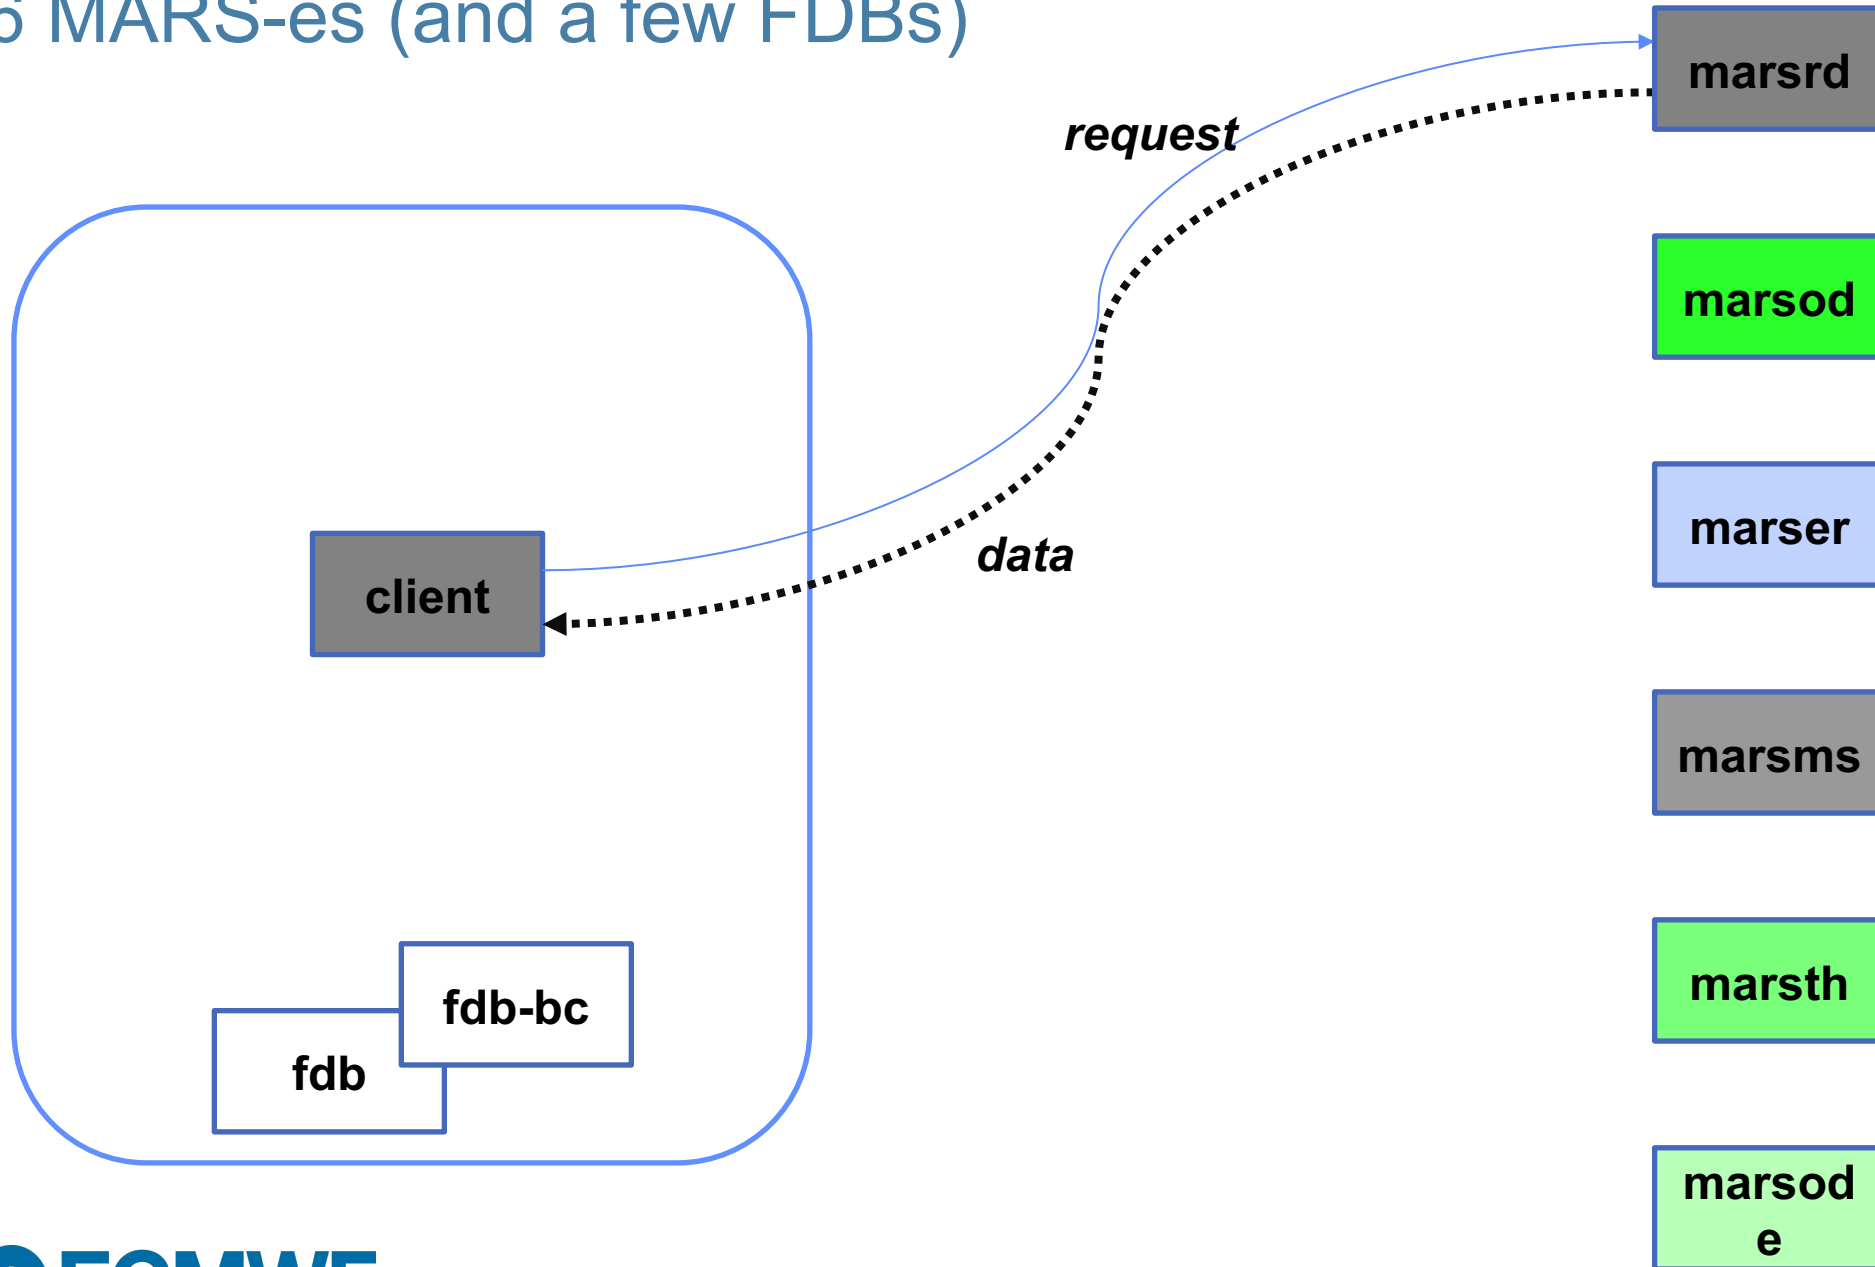
- Client/Server architecture
- Clients: Atos HPCF and Web MARS virtual machine
- Servers: supercomputers, dedicated servers
- Several databases
- Tape library



MARS Architecture – Servers

- Reports Database (RDB), on-line observations (for Operations only)
- Fields Database (FDB)
 - Data produced by most recent cycles or experiments
 - Very fast access (on-line data)
 - Suitable for model input
- ODB server, on-line ODB on supercomputer
- Main Archives (multiple servers)
 - Dedicated Linux servers / clustered architecture
 - Terabytes of disk space
 - Tape management SW: HPSS
 - Automated Tape Libraries: IBM and SpectraLogic

6 MARS-es (and a few FDBs)



Introduction – Some figures

- > 450 PB of data in
- > 400 TB archived per day
- ~ 8000 users with access to MARS
- 1000 active users/day executing ~ 1.5 million requests/day
- ~ 350 TB retrieved daily on average (100 TB / day in 2017) with peaks over 600 TB/day
- Only 2.6 % of the data is on disk
- Operational forecast since 1985
- Analysis, forecast and observations since 1900 (ERA-20C)

What is a PiB?

(Assuming reading from/writing to disk at 100 MiB/s)

	Bytes	Seconds	Days	Months
MiB	1,048,576	0.01		
GiB	1,073,741,824	~10		
TiB	1,099,511,627,776	10,485	0.12	
PiB	1,125,899,906,842,624	10,737,418	124	> 4

Copying 1 PiB |



| 99 days, 13:07:39

remaining

Tape libraries



Terminology – Forecast lead times

- **Medium-range**

the high-resolution and the ensemble forecasts of weather, at the space and time-scales represented by the relevant model, up to 10 and 15 days ahead, respectively, and the associated uncertainty

- **Extended-range (monthly)**

ensembles of individual forecasts and post-processed products of average conditions (e.g. weekly averages) up to 46 days ahead, and the associated uncertainty

- **Long-range (SEAS)**

ensembles of individual forecasts and post-processed products of average conditions (e.g. monthly averages) up to 13 months ahead, and the associated uncertainty

Terminology – ... more

- Re-forecast

forecasts run for past decades necessary to estimate the model climate and the level of skill and to generate some of the operational products

- IFS

‘Integrated Forecasting System’, *the system* used at ECMWF

<https://confluence.ecmwf.int/display/DAC/File+naming+convention+and+format+for+real-time+data>

www.ecmwf.int/en/forecasts/documentation-and-support

Meteorological content – Reanalysis datasets

Browse reanalysis datasets

ECMWF uses its forecast models and data assimilation systems to 'reanalyse' archived observations, creating global data sets describing the recent history of the atmosphere, land surface, and oceans.

Dataset	Time period	Atmosphere	Atmospheric composition	Ocean waves	Ocean sub-surface	Land surface	Sea Ice	Observation Feedback Archive	Download using MARS web interface
ERAS	1940-present	✓		✓		✓			Get ERAS from the Climate Data Store
	1950-1978 preliminary	✓		✓		✓			Get ERAS preliminary from the Climate Data Store
ERAS-Land	1950-present					✓			Get ERAS-Land from the Climate Data Store
ERA-Interim	1979-August 2019	✓		✓		✓			Users are strongly advised to migrate to ERAS from the Climate Data Store (CDS)
ERA-Interim/Land	1979-2010					✓			Users are strongly advised to migrate to ERAS-Land from the Climate Data Store (CDS)
CERA-SAT	2008-2016	✓		✓	✓	✓	✓		
CERA-20C	1901-2010	✓		✓	✓	✓	✓		
ERA-20CM	1900-2010	✓		✓		✓			
ERA-20C	1900-2010	✓		✓		✓	✓		
ERA-20CL	1900-2010					✓			Contact ECMWF
ERA-40	1957-2002	✓		✓		✓			
ORAS4	1958-2015			✓					Get ORAS4 from the Ocean Synthesis/Reanalysis Directory
ORAP5	1979-2013			✓		✓			Get ORAP5 from the Ocean Synthesis/Reanalysis Directory
ORAS5	1975-present			✓		✓			Get ORAS5 from the Ocean Synthesis/Reanalysis Directory
ORA-20C	1900-2010			✓		✓			Get ORA-20C from the Ocean Synthesis/Reanalysis Directory
CAMS RA (EAC4)	2003-June 2022	✓	✓						Get CAMS RA (EAC4) from the Atmosphere Data Store
CAMS GHG RA (EGG4)	2003-2020	✓	✓						Get CAMS GHG RA (EGG4) from the Atmosphere Data Store
MACC	2003-2014	✓	✓						Users are advised to use CAMS global reanalysis (EAC4) from the Atmosphere Data Store (ADS)

Atlas and Picture Gallery
An atlas of key atmospheric variables was produced for ERA-40 and there is picture gallery for ERA-15


ERA5 explorer

To improve our service, we need to hear from you! Please complete [this very short survey](#). Thank you.

Overview Application Documentation Source code

This application provides visualisations of historical climate statistics for any location around the world. Click anywhere on the interactive map or search for a city to explore the typical monthly climate and discover how the climate has changed over the past forty years.

This application is driven by ERA5, the fifth generation ECMWF atmospheric reanalysis of the global climate. ERA5 describes the global history of the atmosphere for the period 1979-2020, using a combination of forecast models and data assimilation systems to 'reanalyse' past observations. As such, the information presented by this application for specific locations are not site-specific observations but rather based on the nearest grid point (nearest 1 degree) to the given location in the ERA5 reanalysis.



User-selectable parameters

- City: the city for which to generate location-specific climate statistics. The performance of the application is maximised for the most populous cities in Europe, along with all European capitals. For lower population and/or non-European cities, you might need to wait up to a minute for the data to be retrieved and processed.
- Variable: global average fields (for the 1981-2010 period) to visualise in the interactive map. The available options are global average temperatures, wind speeds and precipitation totals.
- Statistics: by searching for a city or clicking on the interactive map, you can view a range of climate statistics based on the nearest ERA5 grid point to the selected location:
 - Climatologies (typical monthly averages) of temperature, precipitation, wind speed and wind direction at the given location, averaged over the period 1981-2010.
 - Temperature and precipitations anomalies comparing each year from 1979-2020 to the long-term average period 1981-2010.
 - Monthly and annual percentages of frost days, summer days and tropical nights, averaged over the period 1981-2010.

More details about the products are given in the Documentation section.

INPUT VARIABLES			
Name	Units	Description	Source
10m u-component of wind	m s ⁻¹	ERA5 monthly averaged eastward component of wind at a height of 10m.	ERA5
10m v-component of wind	m s ⁻¹	ERA5 monthly averaged northward component of wind at a height of 10m.	ERA5
Air temperature	K	ERA5 3-hourly temperature of air at 2m above the surface.	ERA5
Instantaneous 10m wind gust	m s ⁻¹	ERA5 monthly averaged wind gust speed at a height of 10m.	ERA5
Mean total precipitation rate	kg m ⁻² day ⁻¹	ERA5 monthly averaged total precipitation rate.	ERA5

OUTPUT VARIABLES			
Name	Units	Description	
Monthly average 10m wind speed and gust speed	m s ⁻¹	The typical monthly average 10m wind speed and gust speed for each calendar month at the given location, averaged over the long-term reference period of 1981-2010.	
Monthly average direction	cardinal direction	The typical monthly average 10m wind direction, derived from the direction of the monthly average u- and v-components of 10m wind.	
Monthly average mean, maximum and minimum temperatures	°C	The typical monthly mean, maximum and minimum temperatures for each calendar month at the selected location.	
Monthly average total precipitation amount	mm	The typical monthly total precipitation amount for each calendar month at the given location, averaged over the long-term reference period of 1981-2010.	
Monthly temperature indices	%	The typical monthly percentage of days which are categorised as frost days (daily minimum temperature < 0°C), summer days (daily maximum temperature > 25°C) and tropical nights (daily minimum temperature > 20°C).	
Yearly precipitation anomaly	%	The difference (as a percentage) between each yearly precipitation total (in mm) for the 1979-2018 period, and the average over the long-term reference period of 1981-2010.	
Yearly temperature anomaly stripes	°C	A visualisation of yearly temperature anomalies, or how much warmer or colder each year in the 1979-2018 period was compared to the long-term reference period of 1981-2010.	
Yearly temperature indices	%	The yearly percentage of days which are categorised as frost days (daily minimum temperature < 0°C), summer days (daily maximum temperature > 25°C) and tropical nights (daily minimum temperature > 20°C).	

Help

[Get help](#)

Licence

Freely available

Publication date

2019-10-01

Related data

[Complete ERA5 global atmospheric reanalysis](#)

[ERA5 hourly data on pressure levels from 1940 to present](#)

[ERA5 hourly data on single levels from 1940 to present](#)

Related applications

[Heating and cooling degree days from 1979 to 2100](#)

[The Universal Thermal Comfort Index \(UTC\) explorer](#)



Access to ECMWF Public DataSets

- TIGGE
- S2S
- Can be download from a web browser or using Web API python client
- <https://apps.ecmwf.int/datasets/>
- <https://confluence.ecmwf.int/display/DAC/Decommissioning+of+ECMWF+Public+Datasets+Service>



Public Datasets

ⓘ The ECMWF Public Datasets service is being decommissioned and access to most datasets **closed** on *June 1st, 2023*. In a final step later this year, access to the remaining multi-model datasets S2S and TIGGE will be migrated to a different system. For more information and alternative access, please visit our dedicated page on the **Decommissioning of ECMWF Public Datasets service**.

Access to these datasets is provided free of charge. Terms and conditions may apply, please check with each individual dataset.

Multi-model

- S2S
- TIGGE

Meteorological content – Data formats

WMO formats

- Fields in GRIB (**GR**id **I**n **B**inary) v 1 and 2, ECMWF local extensions
 - Spherical Harmonics (upper-air fields, T_{CO}1279)
 - Gaussian Grid (surface fields, O1280)
 - Latitude/Longitude (wave and ocean products)
- Observations in BUFR (**B**inary **U**niversal **F**orm **R**epresentation)
 - Instrument specific

ECMWF/IFS format

- ODB (**O**bservational **D**ata **B**ase)
 - Observation feedback

MARS language

Mechanism to *name* fields

Request syntax:

```
verb,  
  keyword1    = value1,  
  ...         = value2,  
  keywordN    = valueN
```

- verb: action to be taken (e.g. retrieve, **list**, read)
- keyword: a known MARS variable, e.g. type or date
- value: value assigned to the keyword, e.g. Analysis or temperature
- Description of all keywords is available in Confluence:

<https://confluence.ecmwf.int/display/UDOC/Keywords+in+MARS+and+Dissemination+requests>

MARS language

- **verb** and **keyword=value** separated by commas, but last one
- Spaces and tab characters are ignored
- *****, **!** and **#** comment until end-of-line
- Directives are not case sensitive
- Values: predefined names, numeric values or strings (filenames)
- Abbreviations: enough letters to uniquely identify keyword or value
- Acronyms: usually initial letters of names
- **/** is used as list separator → specify pathnames in quotes

MARS language – Retrieve request

<code>retrieve,</code>		<code>action</code>
<code>class = od,</code> <code>stream = oper,</code> <code>expver = 1,</code>		identification
<code>date = -3,</code> <code>time = 12,</code>		date & time related
<code>type = analysis,</code> <code>levtype = model levels,</code> <code>levelist = 1/to/137,</code> <code>param = 130,</code>		data related
<code>grid = 2.5/2.5,</code>		post-processing
<code>target = "analysis"</code>		storage

MARS language – Identification of archive

- class** ECMWF classification (od, rd, ea, ...)
- stream** originating forecasting system or (oper, wave, enfo, seas, ...)
- expver** version of the experiment (01 operational, 11, aaaa)
- domain** area covered by the data (Global, Mediterranean, ...)
- origin** originating centre of the data (kwbc, egrr, ...)
- system** seasonal forecast operational system (1, 2, 3)
- method** to specify how the seasonal forecast is produced, e.g. in System 2, method=0 for runs without ocean assimilation (0, 1, ..., 3)

MARS language - Date & time

time	base time or observation time (00, 06, 09:30, ...)
date	base date of the model (-1, 20010225, ...)
step	forecast time-step [hours] from base time (12, 24, 240, ...)
reference	reference forecast time step for EPS tube (96,...)
refdate	date of real-time forecast associated to re-forecast/hindcast (stream=mnfh)
hdate	base date of a re-forecast/hindcast (stream=enfh)
range	observations: period in minutes from base time (360,...) ocean fields: extension of the time series/average
fcmonth	month from seasonal forecast base date (1, 6, ...)
fcperiod	period, in days, for an averaged field (26-32)

MARS language – Fields

type	type of field (an, fc, ...)
levtype	type of level (pl, ml, sfc, pt, pv)
levelist	levels for the specified levtype (off if levtype=sfc)
param	meteorological parameter (t, temperature, 130, 30.128)
number	ensemble member (1, 2, ...)
channel	brightness temperature frequency band
diagnostic, iteration	sensitivity forecast products
frequency, direction	2-d wave spectra products
product, section, latitude, longitude	ocean products

MARS language – Observations & images

type	type of observations or images (ob, fb, ai, af, im)
obstype	observation subtype (s, air) or image channel
ident	WMO observation station number or satellite identifier
duplicates	whether duplicated observations are to be kept or not
block	WMO block number for observation
time	analysis time (types ai, af) or observations time (types ob, fb, im)
range	denotes the period, in minutes, starting from time

MARS language – Storage

- target** UNIX pathname where retrieved data is stored
- source** UNIX pathname from where to read data
- fieldset** temporary storage; can be considered a MARS variable

Unix pathnames (using /) have to be enclosed in quotes, e.g.

```
target = “/scratch/{uid}/analysis”
```


MARS language - Post-processing (1/2)

grid


output grid mesh

- Latitude/longitude increments in degrees (2.5/2.5)
- Type and resolution of Gaussian grid, e.g.
 - grid = F320 – full (or regular) Gaussian grid
 - grid = N320 – ECMWF original reduced Gaussian grid (only selected resolutions supported)
 - grid = O320 – ECMWF octahedral (reduced) Gaussian grid:
<https://confluence.ecmwf.int/display/FCST/Introducing+the+octahedral+reduced+Gaussian+grid>


All above with 320 latitude lines between the pole and equator
F80 regular (or full)



N80 original reduced



O80 octahedral reduced




MARS language - Post-processing (2/2)

area desired sub-area in degrees (north/west/south/east) e.g.
area=75/-25/10/60


frame number of grid points from sub area inwards (5)

**truncation,
intgrid**



rotation lat/lon of South Pole

accuracy number of bits per data value in GRIB (16)



MARS language – Execution control

expect number of expected fields (1000, any, ...)

database where to look for the data

use hint about frequency of use (infrequent)

MARS language – Values

- Expected number of fields is computed by multiplying number of values after expansion of ranges

date = 20020101/to/20020131, 31 days

param = 165.128/166.128/235.128, 3 parameters

93 fields in total

- Certain keywords accept **all** as valid value

levelist = all

- Most keywords accept **off** as valid value

levtype = surface,

levelist = off

Request examples – Operational analysis

Retrieval of sea surface temperature for first 10 days of May 2002, all synoptic times. It retrieves 40 fields.

```
retrieve,  
    class    = od,  
    stream   = oper,  
    expver   = 1,  
    date     = 2023-10-01/to/2023-10-10,  
    time     = 00/06/12/18,  
    type     = an,  
    levtype  = sfc,  
    param    = 34.128,  
    target   = "an.grb"
```

Retrieving data – Calling MARS in a script

- directives from input stream

```
mars <<EOF  
retrieve,  
  type    = an,  
  date    = -7,  
  target  = "$SCRATCH/my_an"  
EOF
```

- directives from file

```
cat > my_request <<EOF  
retrieve,  
  type    = an,  
  date    = -7,  
  target  = "$SCRATCH/my_an"  
EOF  
mars my_request
```

Retrieving data – Hints

- Default values: minimize their use
- No semantic check (only syntax is checked)
- MARS messages
 - INFO request execution and report
 - WARNING unusual aspect of execution
 - ERROR system or data errors
 - FATAL terminates execution

MARS Catalogue – apps.ecmwf.int/mars-catalogue/

Web interface to entire archive content

- Content browsing of *every* field in the archive
 - more up to date than static content documentation
- URL based on MARS requests (can be edited & bookmarked)
- Real-time (dynamic access to metadata)
- Create MARS requests (without checking availability)
- Check availability of data
- Retrieval in GRIB and NetCDF for few fields

MARS Catalogue

- <https://apps.ecmwf.int/mars-catalogue/>



MARS introduction, basic concepts, and advance us

ECMWF

MARS Catalogue Current activity- Reference+

MARS Catalogue

The MARS Catalogue allows you to browse the contents of the whole MARS archive.

Choose the class:

- Operational data
 - Operational archive class=od
- ECMWF Re-Analyses
 - 15 years reanalysis (ERA15) class=er
 - 40 years reanalysis (ERA40) class=e4
 - ERA Interim class=e1
 - ERA-CLIM2 coupled reanalysis of the 20th-century (CERA-20C) class=ep
 - ERA-CLIM2 coupled reanalysis of the satellite era (CERA-SAT) class=et
 - ERA-CLIM model integration for the 20th-century (ERA-20CM) class=em
 - ERA-CLIM reanalysis of the 20th-century using surface observations only (ERA-20C) class=e2
- Copernicus
 - Copernicus Atmosphere Monitoring Service (CAMS, previously MACC) class=mc
 - Copernicus Climate Change Service (C3S) class=c3
 - Copernicus Emergency Management Service (CEMS) class=ce
 - Copernicus Emergency Management Service - Global wildfire class=gw
 - Copernicus regional reanalysis class=rr
 - ERAS class=ea
 - ERAS/LAND class=l5
- ECMWF Experiments
 - Research department class=rd
 - Test class=te
- Special datasets
 - DEMETER class=dm
 - Data Targeting System class=dt
 - ECSN class=cs
 - ELDAS class=e1
 - ENSEMBLES class=en
 - EURO4M class=rm
 - MERSEA class=me
 - NOAA/CIRES 20th Century Reanalysis version II class=nr
 - PROVOST class=pv
 - Sub-seasonal to seasonal prediction project (S2S) class=s2
 - TIGGE class=ti
 - TOST class=to
 - UERRA class=ur
 - WMO Lead Centre Wave Forecast Verification class=lw
 - YOPP class=yp
 - YOTC class=yt
- Member States Projects
 - ALADIN-LAEF class=la
 - COSMO-LEPS class=co
 - Member States projects class=ms
- Member States IFS Experiments
 - Austria class=at
 - Belgium class=be
 - Denmark class=dk

MARS Catalogue

The screenshot shows the ECMWF website homepage. At the top, there is a navigation bar with the ECMWF logo and links for Home, About, Forecasts, Computing, Research, Learning, and Publications. Below the navigation bar is a row of flags representing various countries. The main banner features a map of Europe with the text "Advancing global NWP through international collaboration". To the right of the map, there is a section titled "Lightning flash density, last 6 hour" with a "View all charts" button. Below the banner, there are several featured articles and sections:

- SCIENCE BLOG:** "Heat stress and the European heatwave of 2023" by Claudia Di Napoli, dated 3 October 2023.
- NEWS:** "Week-long high-performance computing event looks into 'Diversifying HPC'" dated 25 October 2023.
- NEWS:** "Autumn 2023 Newsletter published" dated 20 October 2023.
- IN FOCUS:** "Fact sheet: ECMWF and the EU's Copernicus programme" dated 11 September 2023.

Below these articles, there are several sections with images and text:

- Forecast charts and data:** "We provide current forecasts, climate reanalyses and specific datasets. These are available via the web, point-to-point dissemination, data servers and broadcasting." Includes a sidebar with links for Charts, Datasets, Quality of our forecasts, Documentation, and Access to forecasts.
- Datasets:** "Forecasts, analyses, climate re-analyses, reforecasts and multi-model data are available from our archive (MARS) or via dedicated data servers or via point-to-point dissemination."
- Accessing forecasts:** "Researchers, WMO members, commercial users and satellite data providers can gain access to ECMWF forecast products under one of our various license agreement types."
- AI/FS BLOG:** "ECMWF unveils alpha version of new ML model" dated 13 October 2023.

At the bottom, there is a section titled "Serving meteorology" with the text: "International collaboration is fundamental to the work of ECMWF, working with the EU and the meteorological community at large, including the World Meteorological Organization, space agencies, and scientists around the world." Below this text are four sub-sections with images and links:

- Contributing to European Programmes:** Copernicus Climate Change Service (C3S), Copernicus Atmosphere Monitoring Service (CAMS).
- Working with space agencies:** Observations, Satellite data monitoring.
- Supporting WMO programmes:** World Meteorological Centre, Data archives, Severe weather forecasting.



MARS activity – apps.ecmwf.int/mars-activity/

Server activity / MARS queue

- Show system activity
- Monitor your requests
- Learn how the queuing system works
 - Reason for queued requests

ECMWF ☰ Help Bojan Kasic

Current activity Reference

MARS server activity

195 active jobs / 271 total jobs

Sort by: 🔍

My activity
No activity for user usbk

Other activity

1210842581@qq.com	RETRIEVE	TI	SFC	bol-webmars-public-svc-blue-005	12 hours 21 minutes 40 seconds
QUEUED On marsth-core, the total number of Requests from ECMWF Web API with more than 10 tape mounts [source] is limited to 1 34,100 fields, 3.89262 Gbytes online, 36.8454 Gbytes on 14 tapes warning: accessing damaged tages XD024500, XD033400, nodes: hpss mvr000 mvr001, libraries: ibmlib08					
3277409407@qq.com	RETRIEVE	TI	SFC	bol-webmars-public-svc-blue-006	13 hours 1 minute 10 seconds
QUEUED On marsth-core, the total number of Requests from ECMWF Web API with more than 10 tape mounts [source] is limited to 1 600 fields, 26.1764 Mbytes on 17 tapes, nodes: hpss, libraries: ibmlib08					
18008349012@163.com	RETRIEVE	TI	SFC	bol-webmars-public-svc-blue-004	12 hours 54 minutes 9 seconds
QUEUED On marsth-core, the total number of Requests from ECMWF Web API with more than 10 tape mounts [source] is limited to 1 161,200 fields, 1.95335 Gbytes online, 6.39095 Gbytes on 33 tapes, nodes: hpss mvr000 mvr001, libraries: ibmlib04 ibmlib08					
camsprod3583	RETRIEVE	MC	20170901	compute-0003.camsprod.copernicus-atmosphere.eu	34 minutes 25 seconds
240 fields, 61.274 Mbytes on 1 tape, nodes: hpss, libraries: ibmlib07					

Reading retrieved data - tools

Some useful tools

- ecCodes
 - `grib_ls, grib_dump, ...`
- Metview examiners
 - `metview -e <grib|bufr|netcdf|odb> <file>`
- CDO - Climate Data Operators
 - See <https://code.zmaw.de/projects/cdo>

MARS Architecture - Request execution

- 1) Check syntax (MARS language and request syntax)
- 2) Print request to be processed
- 3) Query all Supercomputer' s FDB
- 4) Query main archives (if data not in FDB)
- 5) Transfer data
- 6) Post-processing while transferring (if needed)
- 7) Report on result

Request execution (1/3) – checking MARS language and request syntax

```
mars - INFO - 20231029.192717 - Welcome to MARS
mars - INFO - 20231029.192717 - MARS_HOME=/usr/local/apps/mars/configs/prod
mars - INFO - 20231029.192717 - MARS Client build stamp: 20231005140652
mars - INFO - 20231029.192717 - MARS Client bundle version: 6.33.16.2
mars - INFO - 20231029.192717 - package mars-client version: 6.33.16
mars - INFO - 20231029.192717 - package mir version: 1.18.0
mars - INFO - 20231029.192717 - package odc version: 1.4.6
mars - INFO - 20231029.192717 - package fdb version: 5.11.23
mars - INFO - 20231029.192717 - package metkit version: 1.10.15
mars - INFO - 20231029.192717 - package eckit version: 1.24.4
mars - INFO - 20231029.192717 - package eccodes version: 2.32.0
retrieve,
  class=od,
  date=-7,
  expver=1,
  levtype=sfc,
  param=165.128/166.128,
  step=24,
  stream=oper,
  time=12:00:00,
  type=fc,
  area = europe,          # As the area is not compatible with the grid,
  grid = 0.57/0.57,      # MARS will expand the area to make it fit.
  target = "/ec/res4/scratch/usbk/10u_10v.grb"

mars - INFO - 20231029.192717 - Processing request 1
```

Request execution (2/3) – print request

```
RETRIEVE,  
  CLASS      = OD,  
  TYPE       = FC,  
  STREAM     = OPER,  
  EXPVER     = 0001,  
  REPRES     = GG,  
  LEVTYPE    = SFC,  
  PARAM      = 165.128/166.128,  
  DATE       = 20231022,  
  TIME       = 1200,  
  STEP       = 24,  
  DOMAIN     = G,  
  TARGET     = "/ec/res4/scratch/usbk/10u_10v.grb",  
  RESOL      = AUTO,  
  AREA       = 73.5/-27/33/45,  
  GRID       = 0.57/0.57
```

Request execution (3/3)

```
Query FDB { mars - INFO - 20231029.192717 - Requesting 2 fields
           mars - INFO - 20231029.192717 - FDB home /home/fdbprod
           mars - INFO - 20231029.192717 - FDB home /home/fdbbc
           mars - INFO - 20231029.192717 - Setting SO_SNDBUF to 33554432 (32.00 M)
           mars - INFO - 20231029.192717 - Current value is 32768 (32.00 K)
           mars - INFO - 20231029.192717 - Setting SO_RCVBUF to 33554432 (32.00 M)
           mars - INFO - 20231029.192717 - Current value is 43690 (42.67 K)
Query main archive because the data wasnot found in FDB { mars - INFO - 20231029.192717 - Calling mars on 'marsod-core', local port is 48587
           mars - INFO - 20231029.192717 - Callback at address 10.100.192.202, port 48587
           mars - INFO - 20231029.192717 - Mars client is on ac6-202.bullx (10.100.192.202) 48587
           mars - INFO - 20231029.192717 - Mars server is on 10.112.5.10 42673
           mars - INFO - 20231029.192717 - Server task is 270 [marsod]
           mars - INFO - 20231029.192717 - Request cost: 2 fields, 25.1858 Mbytes online, nodes: mvr004 [marsod]
           mars - INFO - 20231029.192717 - The efficiency of your requests in the last 12 hours is 100% [marsod]
Transferring data and postprocessing { mars - INFO - 20231029.192717 - Transferring 26409176 bytes
           mars - INFO - 20231029.192718 - 2 fields retrieved from 'marsod'
           mars - INFO - 20231029.192718 - 2 fields have been interpolated on 'ac6-202.bullx'
Reporting results { mars - INFO - 20231029.192718 - Request time: wall: 1 sec
           mars - INFO - 20231029.192718 - Read from network: 25.19 Mbyte(s) in < 1 sec [140.77 Mbyte/sec]
           mars - INFO - 20231029.192718 - Writing to target file: 36.09 Kbyte(s) in < 1 sec [26.19 Mbyte/sec]
           mars - INFO - 20231029.192718 - Memory used: 114.19 Mbyte(s)
           mars - INFO - 20231029.192718 - No errors reported
```

Retrieving data

Request scheduling

- Queuing system
 - Priorities: user, request age, request cost (number of tapes and fields)

Data collocation

- MARS tree
- Archive objects (for OD data)
 - 1 file per month of AN (1 level type, all times, levels, params)
 - 1 file per forecast (1 level type, all steps, levels, params)
 - 1 file per EPS (1 level type, all steps, members, levels, params)
 - 1 file per month of ERA 5 FC (1 level type, all levels, times, steps, params)

Retrieving data – Efficiency

- Explore data in archive catalogue - collocation
- Estimate amount of data (list command)
 - Number of fields (up to hundreds of thousand / request)
 - Data size (up to 100 Gigabytes / request)
- Check computing resources: quota, CPU time, ...
- Use local target disk (e.g. \$SCRATCH)
- Retrieve as much data from the same tape as possible
- **Avoid constantly accessing the same tape**
- Do not create unnecessary sub-archives in ECFS

Retrieving data – Data access

- Archived data
 - Available to all registered users
- Current (valid) data, i.e. data for which the value of
$$(\text{DATE} + \text{TIME} + \text{STEP}) + 24 \text{ hours} \geq \text{current date/time}$$
 - Needs special registration
 - Contact your Computing Representative
- Boundary Conditions Project & COSMO-LEPS
 - Restricted to participating MS / individual users
- Restrictions for Observations, TIGGE, EUROSIP...
- Data is available according to dissemination schedule, see
 - <https://confluence.ecmwf.int/display/DAC/Dissemination+schedule>
 - For time-critical retrievals, use time-critical framework (option 1)

Web API

- To access MARS and ECMWF Public Datasets in batch
 - Delivers data directly to the users' machine
- Alternative to retrieve – transfer jobs on the HPC
- Requirements
 - Computer or Web User ID
 - client library, e.g. Python
 - API key

See <https://confluence.ecmwf.int/display/WEBAPI/Access+MARS>

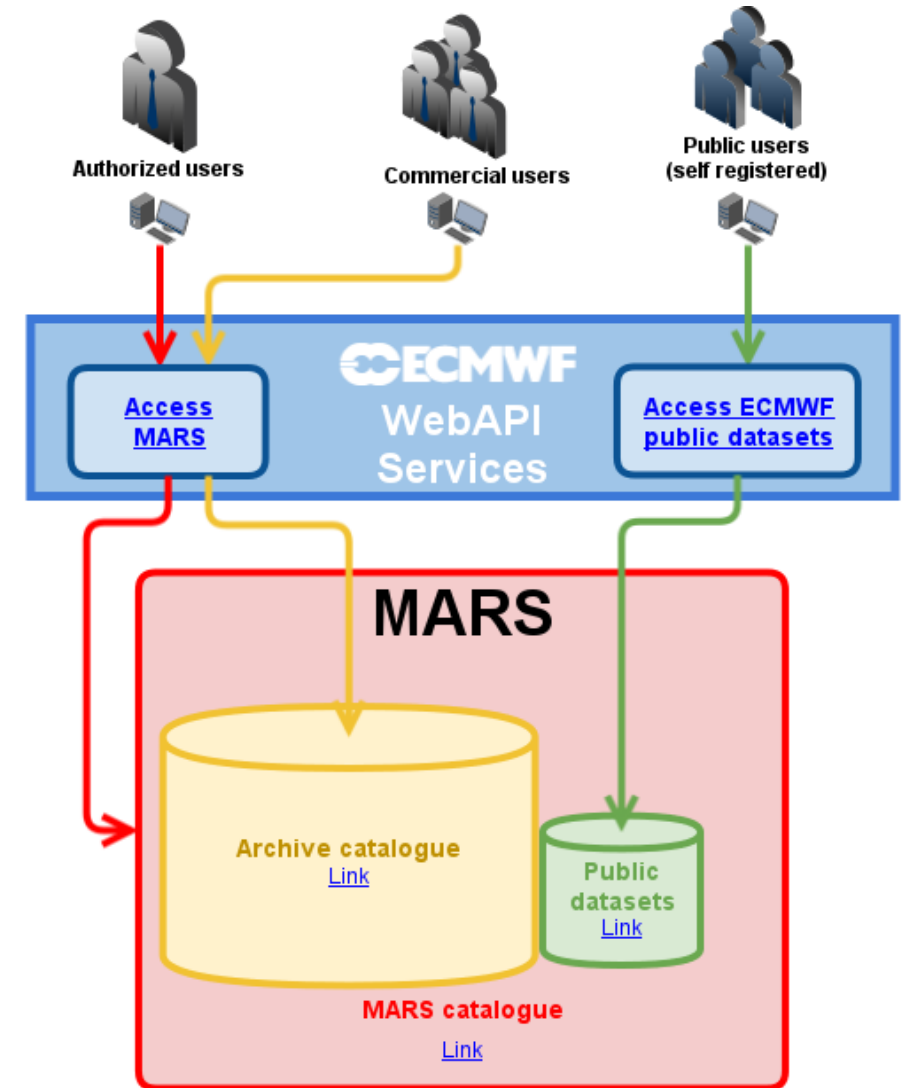
Web API

The Web API is a thin client

- Does not use ecCodes nor any interpolation library
 - therefore any functionality requiring decoding of fields (e.g. compute, read, write, multi-target) is not supported
- Only one request per MARS call is recommended. Only one/two active request allowed currently.
- Target file cannot exceed 75GB and 600,000 fields
- Provides access to full MARS archive to users with the access

Data Server – Web API

Who	Data discovery	Access method
Public user	<u>Public Datasets</u>	<u>Access ECMWF Public Datasets</u>
Commercial user	<u>Archive catalogue</u>	<u>Access MARS</u>
Authorized user	<u>MARS catalogue</u>	<u>Access MARS</u>



<https://confluence.ecmwf.int/display/WEBAPI/ECMWF+Web+API+Home>

Web API Python example

```
#!/usr/bin/env python
from ecmwfapi import ECMWFService

server = ECMWFService("mars")
server.execute(
    {
        "class": "od",
        "date": "20231015",
        "expver": "1",
        "levtype": "sfc",
        "param": "167.128",
        "step": "0/to/240/by/12",
        "stream": "oper",
        "time": "00",
        "type": "fc"
    },
    "target.grib")
```

Examples for retrieving large datasets efficiently:

<https://confluence.ecmwf.int/display/WEBAPI/Retrieval+efficiency>

Additional resources

- MARS documentation
<https://confluence.ecmwf.int/display/UDOC/MARS+user+documentation>
- Web Applications
apps.ecmwf.int/mars-catalogue/
apps.ecmwf.int/mars-activity/
- FAQ
<https://confluence.ecmwf.int/display/UDOC/MARS+FAQ>
- ECMWF real-time datasets
www.ecmwf.int/en/forecasts/datasets
- IFS Documentation
www.ecmwf.int/en/forecasts/documentation-and-support/changes-ecmwf-model/ifs-documentation
- ecCodes Documentation
<https://confluence.ecmwf.int/display/ECC/ecCodes+Home>

MARS Practical

Point your browser to

<https://confluence.ecmwf.int/display/UDOC/MS+computing+training+courses>

or on <https://confluence.ecmwf.int/> navigate to

[User Documentation > MARS user documentation > MS computing training course](#)

and follow the instructions

More examples are available on:

<https://confluence.ecmwf.int/display/UDOC/MARS+example+requests>

if you wish to practice later at home.

MARS – Advanced use

List

- Alternative to the archive catalogue on the web:
 - Amount of data
 - Number of fields
 - Number of tapes. (directive '**output=cost**' required)
 - Suitable for batch mode
- Default is **all**, except for **class**, **expver**, **stream**, **type** and **date**.
- It does list only the archive, not the Fields Database.
- Can keep a report specifying **target**.

List – an example

```
list,  
# output = cost,  
class = od,  
date = 20200101,  
expver = 1,  
levtype = sfc,  
param =  
8.128/22.228/29.228/33.128/45.128/57.128/59.128/134.128/136.128/  
141.128/151.128/164.128/165.128/166.128/167.128/169.128/186.128/  
187.128/188.128/205.128/206.128/246.228/247.228,  
step =  
0/3/6/9/12/15/18/21/24/27/30/33/36/39/42/45/48/51/54/57/60/63/66  
/69/72/75/78/81/84/87/90/93/96/99/102/105/108/111/114/117/120/12  
3/126/129/132/135/138/141/144/150/156/162/168/174/180/186/192/19  
8/204/210/216/222/228/234/240,  
time = 00/12,  
type = fc,  
stream = oper
```

List output

```
class = od
date = 2020-01-01
expver = 1
file[0] =
hpss:/mars/prod/od/o/oper/fc/sfc/marsodoper/0001/fc/20200101/sfc/1200/973003.20200206.0150
29
id = 973003
levtype = sfc
month = 202001
stream = oper
time = 12:00:00
type = fc
year = 2020
file length missing offset param step
0 13204588 . 30376750 8.128 0
0 13204588 . 188831806 22.228 0
0 13204588 . 280684868 29.228 0
0 13204588 . 320298632 33.128 0
...
0 13204588 . 254639830208 205.128 240
0 13204588 . 254653034796 206.128 240
0 13204588 . 255068976704 246.228 240
0 13204588 . 255082181292 247.228 240
```

Grand Total:

=====

```
Entries : 2,990
Total : 40,339,676,520 (37.5693 Gbytes)
mars - INFO - 20231031.155352 - Request performed on database 'marsod'
mars - INFO - 20231031.155352 - Memory used: 31.47 Mbyte(s)
mars - INFO - 20231031.155352 - No errors reported
```

Use directive
'output=cost' for summary
report.

With “output = cost” option

```
size=40339676520;  
number_of_fields=2990;  
online_size=40339676520;  
off_line_size=0;  
number_of_tape_files=0;  
number_of_disk_files=47;  
number_of_online_fields=2990;  
number_of_offline_fields=0;  
number_of_tapes=0;
```

Read: filtering

- Read requests can be used to filter/manipulate already retrieved data.
- Read UNIX file specified by **source**.
- Data written to a file specified by **target** .
- Read doesn't need all directives.

```
read,  
  source      = "myfile",  
  levelist    = 1000,  
  grid        = 2.5/2.5,  
  target      = "only_1000"
```

Not available with
the MARS WebAPI.

Multi-target

- Can organise GRIB target files depending on values of MARS language keywords or of ecCodes key
- MARS Keyword (as echoed by MARS) enclosed in **square** brackets:

```
retrieve,  
  type      = an,  
  expver    = 1,  
  levtype   = sfc,  
  date      = 20010101,  
  time      = 00/06/12/18,  
  target    = "analysis.[time]"  
  ...
```

Not available with
the MARS WebAPI.

```
mars - INFO - 20020515.1237 - Creating target name: analysis.0000  
mars - INFO - 20020515.1237 - Creating target name: analysis.0600  
mars - INFO - 20020515.1237 - Creating target name: analysis.1200  
mars - INFO - 20020515.1237 - Creating target name: analysis.1800
```

Set env. variable **MARS_MULTITARGET_STRICT_FORMAT=1** to use directive values as reported by MARS.

Multi-target

- ecCodes key name enclosed in curly brackets:

```
retrieve,  
  type          = an,  
  param         = t/u/v,  
  expver = 1,  
  levtype = sfc,  
  date          = 20010101,  
  time          = 00/06/12/18,  
  target        = "analysis.{shortName}"  
  ...
```

Not available with
the MARS WebAPI.

mars - INFO - 20140222.164526 - Creating target name: analysis.t

mars - INFO - 20140222.164527 - Creating target name: analysis.u

mars - INFO - 20140222.164527 - Creating target name: analysis.v

Multiple requests

- More than one request in a single call to MARS.
- Append to target:

```
retrieve,  
  date      = 20010101,  
  time      = 12,  
  type      = an,  
  target    = "analysis"
```

```
retrieve,  
  date      = 20000101,  
  time      = 12,  
  type      = an,  
  target    = "analysis"
```

Available with the
MARS WebAPI and
client.

Multiple requests

- Parameter inheritance
 - Parameters not set in the second request (and subsequent ones) are inherited from the previous one:

```
retrieve,  
  class      = od,  
  expver     = 1,  
  stream     = oper,  
  date       = -10,  
  time       = 12,  
  type       = an,  
  target     = "analysis"
```

```
retrieve,  
  type       = fc,  
  step       = 24/to/72/by/24,  
  target     = "forecast"
```

- Unwanted inherited parameters are removed by specifying “**off**”,

Fieldset

- Temporary storage for further processing:

retrieve,

class	= od,
expver	= 1,
stream	= oper,
date	= -10,
levtype	= ml,
levelist	= 1/to/60,
param	= t,
time	= 12,
type	= an,
fieldset	= analysis

Not available with
the MARS WebAPI.

Fieldset

- ‘analysis’ above can be seen as a variable to be referenced in a further request.
- At the end of the call to MARS, all fieldsets are released.
- Write requests save fieldsets into UNIX files:

```
write,  
  fieldset    = analysis,  
  target      = "data.grb"
```

Not available with
the MARS WebAPI.

Compute

- Computations on GRIB fields with same shape:
 - `fieldset`
 - `formula`
 - Scalar values allowed
 - Predefined functions in formula, e.g.

```
compute,  
formula      = "x/2+log(y)*x",  
fieldset     = z
```

where x and y are two fieldsets which have been initialised beforehand.

Not available with
the MARS WebAPI.

Compute

- Mixing fields and scalars:
 - Compute on 2 fields is a field
 - Compute on a field and a scalar is a field
 - Compute on 2 scalars is a scalar
- Bitmaps and missing values:
 - Not considered on computations but copied
- GRIB headers on result of `compute` are incorrect. They are copied from the first fieldset!

Not available with
the MARS WebAPI.

Compute: example

```
retrieve,  
  class = od,  
  expver = 1,  
  stream = oper,  
  type = analysis,  
  date = -10,  
  time = 12,  
  levtype = sfc,  
  param = 165.128, # u  
  grid = 1/1,  
  fieldset = u  
retrieve,  
  param = 166.128, # v  
  fieldset = v  
compute,  
  formula = "sqrt(u*u + v*v)",  
  fieldset = speed write,  
  fieldset = speed,  
  target = "windspeed"
```

Not available with
the MARS WebAPI.

How to ask for help with data download from MARS

- Use: <https://support.ecmwf.int>
- Always provide your MARS request and the full output
- If you are planning to download large amount of data (e.g., forecasts for an entire year or longer period), you can get in touch to make sure the data is download efficiently
- We are monitoring MARS activities. Users with very inefficient activity might be contacted by us to improve the efficiency.

Practicals: Advanced MARS usage

- <https://confluence.ecmwf.int/display/UDOC/MS+computing+training+courses#MScomputingtrainingcourses-AdvancedMARSusage>