

ESO Phase 3 Data Release Description

Data Collection	MUSE
Data Provider	ESO, Quality Control Group
Document Date	2023-12-07
Document version	1.8.3
Document Author	Reinhard Hanuschik/Danuta Dobrzycka

Abstract

This is the release of reduced IFU datacubes from the MUSE¹ spectrograph, taken in the Wide Field Mode. MUSE, the **M**ulti-**U**nit **S**pectroscopic **E**xplorer, is an Integral Field Spectrograph located at the VLT UT4 telescope. It has a modular structure composed of 24 identical IFU modules. The instrument samples almost the full optical wavelength range with a mean resolution of 3000. In the Wide-Field mode (WFM), it has a 1 squared arcmin field of view (FOV) sampling the sky with 0.2 arcseconds spatial pixels. In this mode, observations are carried out under natural seeing (WFM-NOAO) or, since 2017, also assisted by the UT4 AO system (WFM-AO). The Narrow-Field mode (NFM) has a FOV of 7.4"x7.4" and samples with 0.025 arcseconds spatial pixels. It is offered since P103 and is supported by laser tomography, as NFM-AO.

This release is an open stream release. It is complete from the begin of operations, 2014-09-13, until present. Data from the Science Verification periods (NOAO in June and August 2014, AO in August and September 2017) are also included. The data content grows with time as new data are being acquired and processed (approximately with monthly cadence and with a delay of 1 or 2 months).

The data have been reduced with the MUSE pipeline, version `muse-1.4` and higher². The data have most of their instrument signature removed: they have been pre-processed, de-biased, flat-fielded, astrometrically calibrated, sky-subtracted, wavelength-calibrated, and fluxed. Their wavelength scale has been corrected to the barycentric reference system. In the last step they usually have been combined (except for documented cases) and resampled. The sky correction either uses user-defined SKY pointings, or – if not available - the sky is determined from the observations directly, applying a sky model. Errors are propagated throughout the reduction and provided in the final datacube.

The processing is performed by the Quality Control Group at ESO in an automated process. The pipeline processing uses the archived, time-matching, quality-controlled, and certified master calibrations. The reduction process is largely automatic. There is an automatic scoring process for the quality control, and a semi-automatic review and certification process for the data products, focusing on non-zero scores.

The data format follows the ESO science data products standard for datacubes³. Each datacube has a 3D image with two spatial and one spectral coordinates in its first extension. The second extension contains the error information.

The datacubes are usually combined from several input files (typically 2 up to 4, but could be as many as 16 or more), as defined by the observer, as a single OB⁴. Before processing we have reviewed the data acquisition pattern. In some cases we have modified the input data selection, in an attempt to optimize the quality and usefulness of the data product (in a schema called *smart association*). When combination is not possible, or makes no sense, we provide *single* datacubes. See

¹ <http://www.eso.org/sci/facilities/paranal/instruments/muse.html>

² Weilbacher P.M., Streicher O., Palsa R.: "MUSE-DRP: MUSE Data Reduction Pipeline", <http://ascl.net/1610.004> and Weilbacher P.M. et al. [2012SPIE.8451E..0BW](https://doi.org/10.1117/1.5011111)

³ Find it under <http://www.eso.org/sci/observing/phase3.html> under 'Quick links', 'ESO SDP standard'. 'Integral Field Spectroscopy: 3D Data Cubes' defines the structure of the MUSE data products.

⁴ OB = Observing block, a single pointing on the sky and the fundamental unit of the VLT observations.

more under 'Data selection'. For OBs with one or more dedicated SKY observations, we provide a second datacube, processed from the SKY observation(s) by treating them as science pointings. They are on average shallower than the OBJECT datacubes (therefore also called *shallow datacubes* here). They are potentially useful as serendipitous science observations.

This data release offers science-grade data products, with the instrumental signature and the sky background removed. They get resampled only once, at the last step. The spectra come with error estimates and with a list of known features and shortcomings. They are considered to be ready for scientific analysis.

Where applicable we have also processed and combined MUSE datacubes beyond the OB limit, to go as deep as possible in those cases where the observing programme planned to do so. These science products are called MUSE-DEEP datacubes and are described in a separate release description.

Disclaimer. Data have been pipeline-processed with the best available calibration data. However, please note that the adopted reduction strategy may not be optimal for the original scientific purpose of the observations, nor for the scientific goal of the archive user. In particular, the user should critically check the sky subtraction since it is impossible to find a strategy that is always suitable without knowing the science case. There might also be rare cases where the final combination of input data went wrong and caused e.g. double sources.

Release Content

The MUSE release is a stream release. The content of the MUSE release grows with time as new data are being acquired and processed. The data are tagged "MUSE" in the ESO archive user interface⁵. On the same interface you can query for the MUSE-DEEP datacubes⁶.

The first MUSE datacubes have been published on 2016-05-31. The release covers all MUSE data from the two WFM-NOAO (natural seeing) and the WFM-AO Science Verification periods and from start of operations. New data are being added at roughly monthly intervals and with a delay of 1 or 2 months after acquisition, when all master calibrations for the corresponding time interval are certified and available in the archive. Datacubes for the NFM-AO mode have been added since April 2019 (P103).

Relation between MUSE and MUSE-DEEP release. Some of the MUSE programmes have acquired repeated OB visits over more than a year. For the decision about a deep target combination, it is important to have the individual OB-based MUSE datacubes processed (this release) before the deep datacube can be processed (MUSE-DEEP release). Therefore there is a natural sequence (MUSE datacubes come first, then followed by the MUSE-DEEP datacubes), and a natural time delay of roughly one year between the availability of the MUSE and the MUSE-DEEP datacubes.

Data Selection

Data selection is based on the following rules:

- instrument=MUSE;
- observing technique (FITS key DPR.TECH) = IFU;
- instrument mode (INS.MODE) = WFM-NOAO-{E or N} and WFM-AO-{E or N} and NFM-AO-N (there is no E for NFM-AO);
- category (DPR.CATG) = SCIENCE;
- type (DPR.TYPE) = OBJECT and also SKY.

Processing of the MUSE science data has a completion rate of nearly 100%, only very few input raw files are rejected. Data with DPR.TYPE = SKY are processed twice: once with the "sky" recipe

⁵ http://archive.eso.org/wdb/wdb/adp/phase3_spectral/form

⁶ Select or de-select them by the proper choice in the "Multi-OB" field.

`muse_create_sky` used to create the SKY products, and once with the “science” recipe `muse_sciptest` to create a science datacube.

No selection is made on the basis of the observing mode (visitor or service). Data with ‘test programme IDs’ starting with 60. are usually not processed, with the exception of data from the Science Verification runs in 2014 and 2017 which have been added to the data release.

Settings. The wide-field modes without and with AO (INS.MODE=WFM-NOAO and WFM-AO) are both supported, as is the narrow-field mode (NFM-AO). The tag *N* means ‘nominal wavelength range’ (480-930 nm), while *E* means ‘extended’ (465-930 nm). The WFM-AO and NFM-AO modes have a gap without signal between about 580 and 596 nm (*N* range), and 576 and 601 nm (*E* range), respectively, due to laser-induced sodium lines.

Datasets, input files. The standard processing scheme for MUSE datacubes is *template-based*: all input MUSE science files from a single template execution (in many but not all cases identical to a single OB execution) are processed into 1 combined datacube, with a second *shallow* datacube produced from the SKY observation(s), if existing. That second datacube is considered a *serendipitous* science observation, typically shallower than the main pointing (since its exposure time is likely to be shorter), and is offered as a stand-alone science datacube.

Looking closer, there are 3 possible cases for the definition of the set of input files that are used for the datacube creation: the already mentioned **standard** OB case, the **optimized** OB case, and the **complex** OB case (see Table 1).

Table 1. Cases of input dataset definition

Number of OB executions	number of templates per OB	number of datasets	case (see below)
1	1	1 (this is the standard OB case and the most frequent one).	STANDARD
1	N > 1	1 or N, depending on the offsets (“ <i>smart association</i> ”) (this is the complex OB case and an exception).	COMPLEX
> 1 (same night)	...	Can legally happen only when the first execution is graded C; results based on circumstances: - each OB execution gets its standard template-based dataset, or - each OB execution gets its <i>modified</i> dataset, or - all data are combined in one single dataset. (This is the optimized OB case and a not so rare exception.)	OPTIMIZED
1 for OBJECT, plus one for SKY; both concatenated ⁷	1	“ <i>Hidden sky</i> ”: the SKY has its own OB and has the wrong value of DPR.TYPE (OBJECT instead of SKY). If this patterns is recognized, we have transformed the data of the second OB into DPR.TYPE=SKY and then handled them accordingly. (This is another instance of the optimized OB case and a not so rare exception.)	OPTIMIZED

STANDARD case. The baseline design for the processing chain is to take all files of a dataset defined

⁷ Through the same CONTAINER_ID and CONTAINER_TYPE=C.

by the same value of TPL.START⁸. If there are SKY files in that dataset, they are used in the processing chain to calculate the sky background, and to produce a second shallow datacube.

There is however also the case when the acquisition pattern is such that the individual exposures of an OB do not overlap with each other at all, or come e.g. in pairs of two which do not overlap with any other pair. Then the combined datacube would not offer any advantage. We have then decided to split the OB into corresponding individual datasets (Figure 1).




Figure 1. Tile acquisition pattern: there are no overlapping exposures. We deliver N single datacubes (4 in this example).

COMPLEX case. If an OB contains more than one template (in Service Mode only), this is likely designed such in order to have smaller data processing units, or in order to minimize risks of degrading ambient conditions. However, without information about the intention of the PI, we have tried a best guess, based on the observed offset patterns (dither patterns):

- With small dither offsets within a template, but large offsets between subsequent templates, we process the datasets by template (this is then a *mapping* strategy).
- With large offsets only, we produce single datacubes only (this is a *survey* strategy).
- An intermediate case are OBs designed to have a sequence of exposure times (from very short to long) on the same pointing, followed by a set of mapped exposures with the same exposure time⁹. This strategy is chosen by PIs who want to avoid saturation on a central bright target, and also aim at mapping the fainter outer structure, within the same OB. There is clearly no way to combine all such exposures in one single datacube. We have then adopted a strategy to combine all exposures with the same exposure time in a combined cube, and also deliver the ultra-short exposures as single datacubes, with the option to suppress the shortest ones because of high risk of pipeline failure.

The careful checking of the usefulness of combining datacubes vs. delivering single datacubes is also taking into account the size of the MUSE datacubes which might be critical in some situations

OPTIMIZED case, “smart association”. For Service Mode OBs, it occasionally happens that they get executed twice, with their first execution graded C because e.g. one out of 4 input files violated some constraint (because e.g. the seeing degraded). Instead of then repeating the whole OB to bring it to grade A, often only one file is repeated. This would formally result in two datasets, neither of them being optimal. Instead, we have tried to re-shuffle the files in the datasets such that we have one good dataset (good seeing, graded A) and one not so good (graded C), see Table 2. This is likely what any archive user would do in case of interactive selection. We then process both datasets and ingest both datacubes in the archive, to give an archive user the option to evaluate both datasets scientifically from any perspective, not just based on seeing. We have not combined all data in one datacube, however.

We call this approach “smart association”: it goes beyond a rule-based approach and is interactive. We felt that this is a reasonable approach to improve the quality and usefulness of the MUSE datacubes.

Table 2. Smart association

⁸ In many cases, this is corresponding to a single OB execution, but we also give examples for other situations in the following text.

⁹ Such a sequence could be (in terms of exposure times): 0.02, 0.2, 2, 20 sec on the central pointing (bright target), and then a set of 300 sec exposures for the extended structure mapping.

Definition of input dataset :	
Formal	Smart
dataset A: #1,2,3 good; #4 not good → datacube A not optimal because of mixed quality	dataset A': #1,2,3,5 good → datacube A' is optimal
dataset B: #5 good → datacube B not optimal because of incompleteness	dataset B': #4 not good → datacube B' incomplete

The decisions taken in the process of “smart association”, by the operator, are based on:

- the OB, template and file patterns;
- the applied offsets;
- the OB grades;
- the OB comments.

OB grades and comments exist only for Service Mode (SM) and are given by SciOps staff in the QC0 process. For Visitor Mode (VM) OBs, we have not tried anything similar.

“Hidden sky”. There is a special case for the OPTIMIZED input dataset definition: concatenated OBs with one OBJECT OB and one SKY OB. If we discover such pattern, the treatment is as described in the above table. Without this special handling the two OBs would deliver two independent datacubes, each without SKY observation, which is obviously far away from optimal. The key ‘OBJECT’ of the SKY products is then modified to ‘Sky_for_<original_value>’.

Note that all input dataset optimizations within the “smart association” scheme are best-effort and come without guarantee, neither for correctness nor for completeness.

MUSE-DEEP. For those targets that qualify for the MUSE-DEEP combination, the dataset definition has been based on the OB datasets (in a later step). Check the MUSE-DEEP release description for more.

Products. Each of the input datasets can produce 1 or 2 COMBINED datacubes, or *N* SINGLE datacubes (where *N* is the number of input raw files):

Table 3. Types of product datacubes

Contains SKY?	Offsets?	product datacube(s)
NO	SMALL	1 COMBINED*
YES	SMALL	2 COMBINED* (one from OBJECT, one from SKY)
NO	LARGE	N SINGLE
YES	LARGE	N SINGLE, plus 1 COMBINED* (from SKY)

*If only one input file exists, then there is only a SINGLE datacube product.

The distinction between SMALL and LARGE offsets is made empirically. Effectively datasets have been split to deliver SINGLE products only if there was no pixel overlap. There are four different types of datacube products, as listed in Table 4.

General processing pattern. An important condition for the MUSE science processing is the existence of master calibrations. Most master calibrations are taken at daily frequency. Twilight flats are taken in about weekly cadence.

Table 4. Types of MUSE datacubes

Product type	from input file	PRO.CATG	occurrence
COMBINED	OBJECT	DATA_CUBE_COMBINED	often
SINGLE	OBJECT	DATA_CUBE_SINGLE	rare
COMBINED	SKY	DATA_CUBE_SKY_COMB	rare
SINGLE	SKY	DATA_CUBE_SKY	often

They are all processed and quality-checked by the Quality Control Group at ESO, close to the date of acquisition, in order to provide quality feedback to the Observatory. All master calibrations used here have been certified, meaning they have been checked for quality and for proper registration

of instrument effects.

De-selection. Data with `OBS.PROG.ID` starting with '60' or '060' have not been processed (observatory programmes for testing), with the exception of Science Verification data taken in June and August 2014 which are being added to the release. Data taken at daytime (with obviously wrong 'SCIENCE' tag) have been ignored.

Data with extremely short exposure times (≤ 10 sec) have not been processed, because they would require a customized sky correction. Such data are sometimes acquired if a field contains a very bright source, and several short exposures are taken to handle saturation issues.

Quality issues for the first Science Verification period. In the first Science Verification period we have encountered quality issues with the calibrations. In particular for dates 2014 June 23, 24, 25, and 27, the tracing and illumination calibrations did not match well the prevailing conditions of the science data, due to particularly low ambient temperature. Find more in the section **Known features and issues**.

Release Notes

Pipeline Description

Find the detailed description of the recipes in the Pipeline User Manual¹⁰, section 9 (recipe reference). Find the pipeline version used for processing in the header of the product datacube, under "HIERARCH ESO PRO REC1 PIPE ID". The version for the initial dataset was `muse_1.4`. Information about the MUSE pipeline (including downloads, manuals, cookbook) can also be found under that URL. The MUSE pipeline has been written mainly by Peter Weilbacher (see Weilbacher et al. 2012 <http://adsabs.harvard.edu/abs/2012SPIE.8451E..0BW> for a description, and Weilbacher et al. 2016 <http://ascl.net/1610.004> for the code reference).

The QC pages¹¹ contain further information about the MUSE data, their reduction and the pipeline recipes for *calibration data*. Monitoring of MUSE performance and quality parameters is provided under the Health Check monitor¹².

Data Reduction and Calibration

Reduction steps, overview. The data reduction uses a cascaded recipe scheme, with two main parts. It is the same for NOAO and AO data. AO data are reduced with the proper AO calibration data.

The first part works on individual input raw files. No combination is done at that stage. First, every input raw file (OBJECT or SKY) is pre-processed with the recipe `muse_scibasic`. Then, the SKY product files (if any) are further processed with the recipe `muse_create_sky` to create the SKY_LINES and SKY_CONTINUUM files for the later sky subtraction. The sky contribution is evaluated by considering the information on the instrument line spread function, which is contained in the LSF_PROFILES master calibration file. Next, the OBJECT product files are processed with the recipe `muse_scipost`, using the SKY products (if existing) for the sky subtraction. At that stage, the SKY products are processed for a second time, this time with `muse_scipost`.

¹⁰ Under the MUSE link in <http://www.eso.org/sci/software/pipelines/>.

¹¹ http://www.eso.org/qc/MUSE/pipeline/pipe_gen.html

¹² <http://www.eso.org/HC>, select MUSE.

After the `muse_scipost` step, all input raw files (OBJECT or SKY) have a SINGLE DATACUBE product, an IMAGE_FOV product (a 2D collapse useful as preview, see below), and a PIXEL_TABLE with the same pixel values as the datacubes, but with the pixel coordinates stored explicitly in a table. From these three products, the IMAGE_FOV and the DATACUBE is kept only if there is no further processing possible or needed (case of SINGLE datacube product, see Table 4).

In the second part of the science cascade, those PIXEL_TABLEs which logically belong together (defining the dataset) are combined. Two steps are necessary: first, the input IMAGE_FOVs (the collapsed products of the first part) are processed with `muse_exp_align` in order to measure the relative alignment of the input data (in order to detect and correct for possible alignment errors due to instrument wobble, see below)¹³. Then, finally, the input PIXEL_TABLEs are processed with `muse_exp_combine` which applies the alignment correction (if any), and resamples the overlapping pixels in order to go deep. It is only at that last step that the input data are resampled. The output of that last step is the COMBINED DATACUBE, and the combined IMAGE_FOV.

Find the overview of the recipes in Table 5.

Table 5. Overview of MUSE science reduction cascade.

Recipe	Number in the figures	Applied to
<code>muse_scibasic</code>	1	single OBJECT or SKY
<code>muse_scipost</code>	2	single OBJECT; SKY treated as OBJECT
<code>muse_create_sky</code>	2a	SKY
<code>muse_exp_align</code>	3	set of N>1 input data
<code>muse_exp_combine</code>	4	set of N>1 input data

Reduction steps, details. In the following we list more details about the reduction cascade.

Part 1, single datacube.

1.1 The main reduction steps for `muse_scibasic` are:

- trim the raw data,
- record overscan statistics,
- subtract the master-bias,
- divide by master-flat-field,
- convert to pixel-table (using input tables like geometry table, trace table, and wavelength calibration table),
- use the provided ILLUM correction frame for the relative correction of illumination between all slices of an IFU,
- divide by normalized twilight cube in order to provide the twilight correction.

The output is the pre-reduced pixel table. Input for the recipe are either an OBJECT or a SKY raw file. For the MUSE data products, all pipeline parameters for this recipe are set to their defaults, except for `nifu=-1` (all IFUs processed).

1.2 The SKY data are then processed with the pipeline recipe `muse_create_sky`, in order to extract the sky continuum and the sky emission line spectrum from an almost empty, PI-defined sky region close to the OBJECT observations. All pipeline parameters are set to default. The output products (SKY_SPECTRUM, SKY_LINES,

¹³ We have decided to apply that `muse_exp_align` step only for up to N=4 input files. For larger datasets the alignment algorithm is likely to have issues and is therefore not used.

SKY_CONTINUUM, SKY_MASK) are fed as calibration products into the next step, `muse_scipost`.

1.3 After pre-processing, the products of `muse_scibasic` are processed further with the recipe `muse_scipost`, still at the single-frame level. If one or more SKY observations exist in the input dataset, their products (from the execution of `muse_create_sky`) are associated to the closest-in-time matching OBJECT file.

The main reduction steps for `muse_scipost` are:

- correct the exposure for differential atmospheric refraction;
- execute the flux calibration;
- execute the correction for telluric absorption;
- subtract the sky contribution: a model for the sky contribution, both continuum and emission lines, is evaluated for each slice and each IFU and then subtracted from the data; it is created by using the information on the instrumental line spread function (stored in the master calibration file `LSF_PROFILES`) and the `SKY_CONTINUUM` and `SKY_LINES` produced from the SKY observations previously processed with `muse_create_sky`, if available; if SKY observations are not available, the `SKY_CONTINUUM` and `SKY_LINES` are directly created from the darkest regions of the field of view;
- provide the barycentric RV correction;
- apply the astrometric solution;
- finally, resample the data into a datacube.

At that stage the resampling is done just “in case”, to have a single datacube product, if at a later stage issues with the combined datacube (the intended final product) should be discovered. The main product for further processing is the pixel table. If the dataset contains only one input file, then the datacube is the final product, and no pixel table is stored.

The pipeline parameters for this recipe are set to their default values, except for

- `skymethod=model` if no SKY file is available from the input dataset (if it *is* available, its products are taken and the default value ‘`subtract_model`’ is used for this parameter),
- `skymodel_fraction=0.2` if no SKY file is available from the input dataset.

1.4 The SKY data, after the initial `muse_scibasic` step, are not only processed with `muse_create_sky` to obtain the sky products, but also with `muse_scipost` like the OBJECT data, in order to produce a *shallow datacube*. In that step, the processing parameters are:

- `skymethod=model`,
- `skymodel_fraction=0.5`.

For a discussion of the `skymodel_fraction` parameter see section “Issues”. The choice of the value 0.5 for SKY data is based on the reasonable assumption that user-defined SKY regions are likely to contain only few sources, if any.

For single datacubes, this is the end of processing. See Figure 2.

Part 2, combined datacube.

2.1 In the second part the pipeline recipes work on the products (pixel tables and FOV images) from all input files together. The recipe `muse_exp_align` is used to create a coordinate offset table for automatic exposure alignment. The alignment

procedure is in principle useful to detect and correct for instrument wobble, but has its own limitations and risks (see section “Issues” for a discussion).

Step 2.1 is skipped if

- number of input files is larger than 4 (this is the acquisition pattern MAP with little or no overlap so that no numerical alignment is possible);
- in cases where at the certification step a misalignment in the combined datacube is seen, and the datacube could be reprocessed without alignment.

2.2 Finally the output OFFSET_LIST table from `muse_exp_align`, and the pixel tables from the individual `muse_scipost` products are combined into the final combined datacube. The steps are:

- weight input reduced pixel tables by exposure time,
- apply alignment values from alignment table,
- resample pixel tables into combined datacube.

Products.

If the input dataset contains 1 OBJECT file, the products are one SINGLE datacube and the corresponding FOV 2D image (Figure 2). This is a relatively rare situation, most datasets contain more than one input file. The same applies for the processing of a single SKY file into a shallow datacube.

If the input dataset contains N OBJECT files and no SKY file, the products are one COMBINED datacube and one corresponding FOV 2D image (Figure 3).

If the input dataset contains N OBJECT files and 1 SKY file, the products are one COMBINED datacube from the OBJECT files and the corresponding combined FOV 2D image, plus one SINGLE datacube from the SKY file and the corresponding FOV image (Figure 4).

If the input dataset contains N OBJECT files and M SKY files, the products are one COMBINED datacube from the OBJECT files and the corresponding FOV 2D image, plus one COMBINED datacube from the SKY files and the corresponding FOV 2D image (Figure 5).

In the following figures, we use the following numbering scheme for the recipes:

<code>muse_scibasic</code>	1
<code>muse_scipost</code>	2
<code>muse_create_sky</code>	2a
<code>muse_exp_align</code>	3
<code>muse_exp_combine</code>	4

All steps executed by the pipeline, in the complex reduction cascade, are stored in the text file that is delivered with each datacube. While that information is technical, it might help for the understanding of the individual steps and might also serve as reference in case a user wants to redo certain reduction steps.




Figure 2. Simplest case: single OBJECT file, SINGLE datacube. This is also the result for a single SKY file. The main product is one SINGLE datacube.

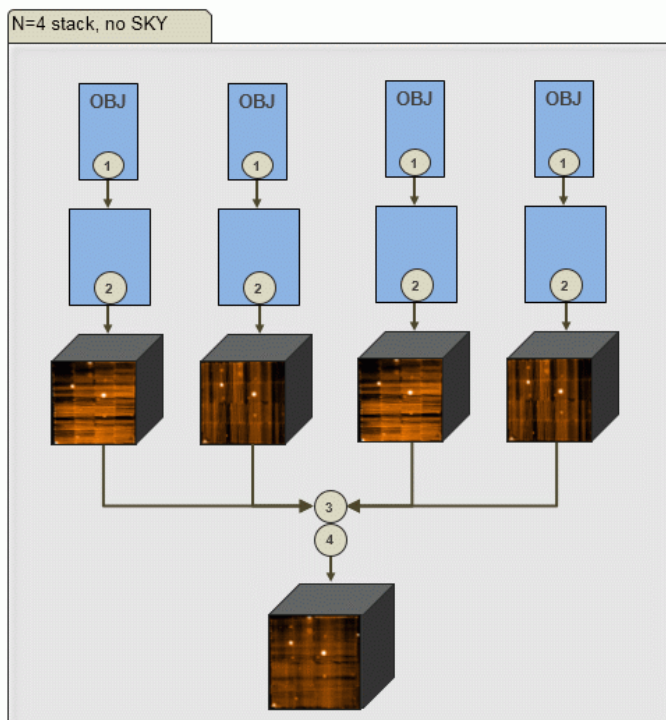


Figure 3. Reduction cascade for N=4 input files, no SKY. Main product is one COMBINED datacube.




Figure 4. Like Figure 3, with 1 SKY file. Main products are one COMBINED datacube as defined by the OBJECT files, plus one SINGLE shallow datacube from the SKY observation.

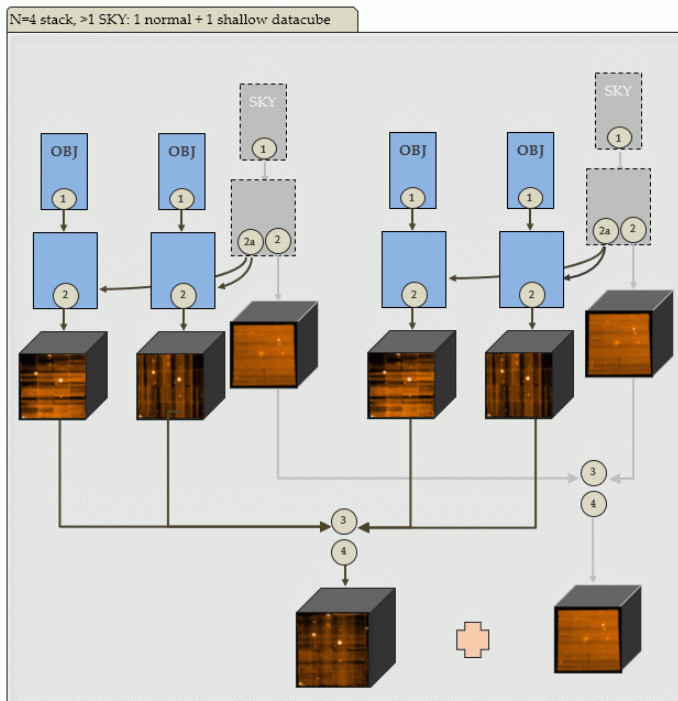


Figure 5. Like Figure 4, with more than 1 SKY file. Main products are one COMBINED datacube as defined by the OBJECT files, plus one COMBINED shallow datacube from the SKY observations.

Master Calibrations used for data reduction¹⁴

Table 6. Set of master calibrations used for data reduction

Type (pro.catg)	name (first part)	content
always required:		
MASTER_BIAS	MU_MBIA	created from 11 raw bias frames; removes bias level and bias structure.
MASTER_FLAT, TRACE_TABLE	MU_MFLA, MU_PFTT	created from 11 raw flat frames; removes gain (fixed pattern) “noise”; traces edges of slices; locates bright and dark pixels
WAVECAL_TABLE	MU_PWCT	created from at least 3 (up to 15) arclamp frames; stores the wavelength solution
TWILIGHT_CUBE	MU_MTWC	created from at least 3 skyflat frames; used to correct relative throughput
STD_RESPONSE	MU_MSDR	created from a STD star observation, used for flux calibration
STD_TELLURIC	MU_MSDT	same, used for correcting telluric absorption lines
provided if SKY observation available in the input dataset:		
SKY_LINES, SKY_CONTINUUM		created by muse_create_sky from input SKY file
additional raw frame used for reduction:		
ILLUM_FLAT		raw frame taken every 2 hrs at the telescope in order to provide an up-to-date measurement of illumination corrections; applied by the pipeline directly as raw frame (without individual processing)
static calibrations:		
GEOMETRY_TABLE	MU_GEOT	static table describing the exact geometry of the location of each IFU slice in the FOV
LSF_PROFILE		static table with the line-spread function
ASTROMETRY_WCS		processed from an observation of an astrometric standard field
EXTINCT_TABLE	MU_GEXT	table with extinction curve

Wavelength scale. The MUSE IFU products are wavelength calibrated. The wavelength scale is *bar-centric*.

Telluric absorption. Telluric absorption lines have been corrected with the STD_TELLURIC file that was derived from a standard star observation (the same as for the flux calibration). The user needs to be aware that this star is observed at the beginning of the night, therefore the conditions of airmass and time proximity for an optimal telluric correction might not be matched. User provided observations of telluric standard stars obtained with the observing template `MUSE_wfm_cal_telluric` or `MUSE_nfm_cal_telluric` and having `DPR.TYPE= STD, TELLU` are not processed in our scheme, as they might not be in the pipeline database of standard stars.

Flux calibration. The MUSE datacubes are flux calibrated, by using the STD_RESPONSE file that was derived from a standard star observation. The MUSE calibration plan foresees the measurement of a flux standard star every night when science is taken. Many OBs however are not requiring strictly photometric conditions. The flux calibration of the MUSE data is therefore not “absolute” in general¹⁵. There is no guarantee that the photometric conditions during the standard star observation are identical to the ones during science acquisition, unless good photometric

¹⁴ Please check the Pipeline User Manual for a description of the calibration data.

¹⁵ Despite the value of the keyword 'FLUXCAL'.

conditions were requested and the OB got graded ‘A’. Likewise there is no guarantee that the photometric conditions during the raw file acquisition (which could have taken up to one hour) were stable. It is also possible that the photometric conditions changed between the OBJECT and the SKY acquisitions. For data taken in Service Mode, the user may want to check the headers for the keywords OB_GRADE (the OB grade in the ESO system) and OB_COMM<n> (multi-line comments about the OB execution given by on-site staff).

Master calibration names and recipe parameters used for reduction. The product header contains a list of all used master calibrations. Look for keys “HIERARCH ESO PRO REC3 CAL<n> NAME” and “... CATG”, with the index n. The used pipeline parameters and their values are listed as “HIERARCH ESO PRO REC3 PARAM<n> NAME” and “... VALUE”. The tag REC3 refers to the final recipe `muse_exp_combine`.

There are also similar header keys inherited from the previous reduction steps, REC2 refers to `muse_scipost`, and REC1 to `muse_scibasic`.

There is also the input raw file of type ILLUM_FLAT which is a single flat-field calibration taken every two hours during the night. It is used within `muse_scibasic` for a differential illumination correction of the individual slices in the raw file.

Products

The final MUSE science data product has two 3D image extensions:

- 3D datacube with 2 spatial dimensions and 1 wavelength axis, with flux-calibrated spatial pixels;
- 3D datacube with the errors.

The following additional FITS file is delivered associated to the MUSE datacube:

- 2D white-light image from collapsed datacube, called IMAGE_FOV.

It is useful for previewing the product file in image viewers like `rtcd`.

In addition, there is an associated text file delivered that contains the combined pipeline log with all executions steps for all participating input files for a given datacube.

Also, there is a set of png files that serve both as QC plot and as preview of the FOV. There is always one for the final datacube. If this datacube is a combined datacube, there are also the corresponding QC plots for each individual datacube¹⁶.

The spectra contain some header keywords added that are related to OB information and to the QC process. They are listed in Table 7.

Table 7. FITS keywords added

Parameter	Values	Meaning
<i>OB related information:</i>		
SM_VM	SM or VM	Data taken in Service Mode or Visitor Mode; VM data

¹⁶ Remember that these individual datacubes are NOT delivered as standard data products but are available through a [special download channel](#).

Parameter	Values	Meaning
		are less constrained in terms of OB properties; they have no user constraints defined and therefore no OB grades.
OB_GRADE	A/B/C/D; X	Immediate grade given by night astronomer, considering ambient conditions checked against user constraints (VM data are formally graded X meaning 'unknown').
OB_COMM<n>	Free text	Any optional comments added by the night astronomer, together with the approximate UT hh:mm (truncated after 200 characters).
<i>QC related information:</i>		
QCFLAG	e.g. 00000010000	QC flag composed of 9 or 11 bits, see Table 8.
QC_COMM<n>	Free text	Comment about the acquisition pattern; comments about quality issues might also exist)

Data Quality

Master calibrations. All used master calibrations have been quality-reviewed and certified at the time of acquisition, as part of the closed QC loop with the Observatory which also includes trending. As part of the certification process for calibration data there is a scoring process to bring non-compliant behavior of the calibrations to the attention of the QC scientist. All these cases have been handled as part of the certification procedure. Hence there is reasonable evidence that the master calibrations catch all instrument properties, as relevant for the reduction, correctly and completely.

QC, review and certification process. The MUSE datacubes have been reviewed and certified by a process involving both automatic scoring and human-supervised certification. Both the single datacubes (output of `muse_scipost`) and the combined datacubes (output of `muse_exp_combine`) are exposed to the QC process.

For single datacubes the QC system scores key parameters like

- NAXIS1/2/3 (the size of the datacube axes; anomalies indicate processing failures);
- NUM_SAT (number of saturated pixels in the raw file);
- maximum correction of wavelength scale by the `muse_scibasic` recipe;
- association quality (proximity of arclamp calibration, completeness of dataset).

For the combined datacubes, the QC parameters are

- differential offset applied by the alignment procedure;
- association quality (same as above) and processing completeness.

The measured values are compared to reference values and scored. A non-zero score flags a potential issue. Any non-zero score is reviewed. A fraction of 20% zero-score datacubes is reviewed as well, selected at random. The combined datacubes are always reviewed, to discover alignment issues visually. Upon review, QC comments can be entered. They are intended to be informative, but they don't claim completeness¹⁷. The QC comments are propagated to the datacube headers.

QC flag. While the QC scheme is an internal process, there is also a scheme to propagate quality information to the user. The header key "QCFLAG" in the MUSE products contains a quality flag that is composed of 11 binary bits in case of combined datacubes, and 9 bits for single datacubes (Table 8). For each bit, the value 0 means "OK, no concern".

¹⁷ For example, the comment "negative background" is often given. But if there is no such comment, this does not necessarily mean "no issue". Check the QCFLAG key for more complete QC feedback.

The flag #2 indicates whether the datacube is from the OBJECT files (value is 0) or from the SKY file(s) (value is 1). SKY datacubes are valid measurements of science targets but have typically shorter exposure times.

Flag #3, if 0, indicates that the dataset includes a user-defined SKY pointing. Usually the sky determination then works best because of the realistic assumption that only weak sources exist in the field, and the algorithm can reasonably well remove them and apply the sky products to the OBJECT pointing. The fluxes in the datacube are then well determined no matter if the field is crowded or has complex background. If flag #3 is 1, the sky has been derived in the OBJECT field itself. The algorithm takes all pixels between the 5% and 20% lowest values in the histogram for the sky determination. This can be a reasonable strategy but for crowded fields, or for fields with complex background emission or extended objects the resulting sky might be systematically too high, which leads to an over-subtracted background and negative fluxes. It is impossible to avoid this in absence of knowledge about the science target. If flag #3 is 1, and flag #8 indicates negative fluxes, the sky of the datacube is likely over-subtracted.

Flag #5 indicates if no twilight flat has been used in the reduction. This can happen rarely and only if no valid (certified) twilight flat calibration was available within +/- 14 days around the science file. Then the relative flux variations between the individual IFUs are stronger than with twilight flat.

Flag #8 is important. As discussed further in the section “Issues” it flags cases when the sky background is likely over-subtracted.

Flag #9 marks possible inaccuracies of the wavelength scale. If there was an issue with sky correction, the differential correction using sky lines might be larger than normal. This happens mostly in cases when the sky lines are underexposed, i.e. for very short exposure times. Therefore, data with extremely short exposure times are treated differently. Until P103 (September 2019), data with exposure times ≤ 10 sec were not processed at all. As of P104 (since October 2019), the WFM data with exposure times ≤ 10 sec, and the NFM data with exposure times ≤ 150 sec, are processed without differential correction using sky lines (parameter --skylines=NONE in the recipe muse_scibasic).

Flag #10 refers to the alignment of the input files and marks suspicious cases where the alignment correction should be double-checked by eye. This could either be cases when no alignment correction was applied (more than 4 input files, or suppression of the alignment recipe), or where the correction was higher than usual.

The final QC flag is composed of all components. 0000000000 marks a COMBINED datacube with extra SKY, no quality issues, from a standard single template. 0010000000 marks a COMBINED datacube without extra SKY but no quality issues. The flag 0110000000 stands for a combined SKY datacube without issues. The flag 00100001000 marks a COMBINED datacube with no extra SKY and symptoms of SKY over-subtraction (crowded field). The flag 00000000010 marks a COMBINED datacube with either no alignment correction at all, or with correction values larger than normal. This does not necessarily imply an issue, but the collapsed white-light file and the QC plots should be inspected carefully.

For single datacubes, the same schema applies, with the two last bits being omitted. For instance, 001000010 flags a single datacube without extra SKY, background negative (rather frequent case for crowded fields). A single datacube with saturation issues is marked as 000001000.

QC plots and previews. The QC and preview plots have been originally developed as quick-look plots for the process quality control. It was felt that they might also be useful to the archive user. They are delivered as associated files along with the products. There are two types of plots:

1. the QC plot for the combined datacube (Figure 6, Figure 7);
2. the QC plot for each single datacube (the datacube itself is usually not delivered) (Figure 8).

Process quality control. The quality of the data reduction is monitored with quality control (QC) parameters, which are stored in a database. The database is publicly accessible and has a browser and a plotter interface¹⁸.

Error propagation. The MUSE datacube has a second extension with the variance of each spatial pixel, in units of $(10^{-20} \text{ erg/s/cm}^2/\text{Angstrom})^2$, as obtained from error propagation through the various steps of the data reduction.

The variance for each pixel in a raw frame is initially computed from the photon noise, the read-out noise and the error of the read-out noise estimate¹⁹. This assumes that the errors of the individual pixels are independent to first order. Then, Gaussian error propagation for each pixel is adopted in each further step of the data reduction cascade.

Limiting magnitude. Each datacube has a QC parameter ABMAGLIM. According to the ESO Science Data Products standard for datacubes it defines the “5-sigma AB magnitude limit for point sources” and measures the noise level of the datacube outside of astronomical sources. It is based on the IMAGE_FOV histogram and assumes a Gaussian distribution around the maximum which is then fitted. The width of the fit, interpreted as Gaussian noise, gives a limiting magnitude. This works generally fine for a FOV that is dominated by background noise and with good background subtraction (extra SKY pointing, or good sky modelling). It might give wrong impressions e.g. in cases with over-subtracted sky, or with extended sources, or with complex sources.

¹⁸ Browser: http://archive.eso.org/qc1/qc1.cgi?action=qc1_browse_table&table=muse_sci_combined

Plotter: http://archive.eso.org/qc1/qc1.cgi?action=qc1_plot_table&table=muse_sci_combined

¹⁹ Gössl & Riffeser 2002 A&A 381, 1095

Bit	Content (if YES, value is 0, otherwise 1)	Motivation
#1 – master sky line fit	no pipeline error upon master sky fit?	catch a pipeline error upon master sky fit (“master sky fit failed with error code 21: the iterative process did not converge.”)
#2 – OBJECT vs. SKY	this datacube comes from OBJECT frames?	if from SKY, this is a shallow datacube, with serendipitous pointing (any targets in the field are not the primary goal of the observation)
#3 – SKY observation	a dedicated (user-defined) sky observation exists?	usually the SKY subtraction is better/more reliable if a dedicated SKY pointing exists (for SKY datacubes this is always 1)
#4 – arc calibration	time difference within 1.5d (previous/this/next night)?	usually daytime calibrations come within 0.5 days after the science observation; if more than a day difference, probability for a mismatch is higher, affecting the wavelength scale error (very rarely violated)
#5 – SKY_FLAT	existing?	flag cases without sky-flat correction (larger error for illumination correction)
#6 – saturated pixels	number of saturated pixels in all input raw frames lower than 300?	flag cases with partial saturation (which cannot be directly discovered in the product datacube)
#7 – number of sources	number of sources found by the pipeline >0?	The source detection algorithm used by the pipeline is quite restrictive, hence a value of 0 does not mean anything but an indication that the alignment procedure might have failed or deliver unreliable results.
#8 – sky subtraction quality	HISTO_17 parameter >-20? (see Figure 6)	Quality of sky subtraction: for crowded fields the sky might get oversubtracted; check for this if this bit is 1.
#9 – wavelength scale quality	LSHIFT_MAX	Quality of wavelength scale: maximum of residual correction done on sky lines, in Angstrom; should be <0.5 A (<0.2 A before October 2019).
The next 2 flags apply to COMBINED datacubes only:		
#10 – alignment	differential offset between individual observations <6e-5 deg (0.2arcs) *AND* all input frames matched?	if offset larger than threshold, this might indicate an alignment issue, please check the IMAGE_FOV carefully for elongated or double images; value is also set to 1 if no alignment correction has been applied
#11 – single template	input dataset from single template?	if no (value=1), we have combined multiple templates (within the same OB) with a somewhat increased probability that ambient conditions changed. Still, this strategy might be superior to have individual datacubes, hence this score parameter is designed just to make the user aware.

Table 8. Definition of QC flags. Find the up-to-date list under the URL

http://www.eso.org/qc/PHOENIX/MUSE/score_bits.txt.




Figure 6. Main QC plot of the combined datacube, featuring: the preview (display of the IMAGE_FOV file); the raw histogram of the 1st input raw frame (close-up of the range 50,000-65,000 ADU); two product histograms (one as a close-up of low fluxes around zero, to check the background subtraction; the other a histogram for the entire dynamic range of the datacube). **At bottom:** a set of QC parameters applicable to the product (Texptime = total exposure time of the datacube, Nsources = number of pipeline-detected sources, as marked on the display; mean_FWHM = median FWHM of point sources, for AO mode; ABMAG_limit = limiting magnitude (depth) of the datacube; delta_t = time difference in days between arclamp calibration and science files; Ninput = number of input OBJECT files; histo mode = flux value for the maximum in the product histogram, also marked by the broken line; histo-1.7 = flux value where histogram value is fallen off by -1.7dex as compared to the mode; score_bit = QC flag as also stored in the header, see Table 8. **On top:** some keywords read from the product file header, like OB name and target.

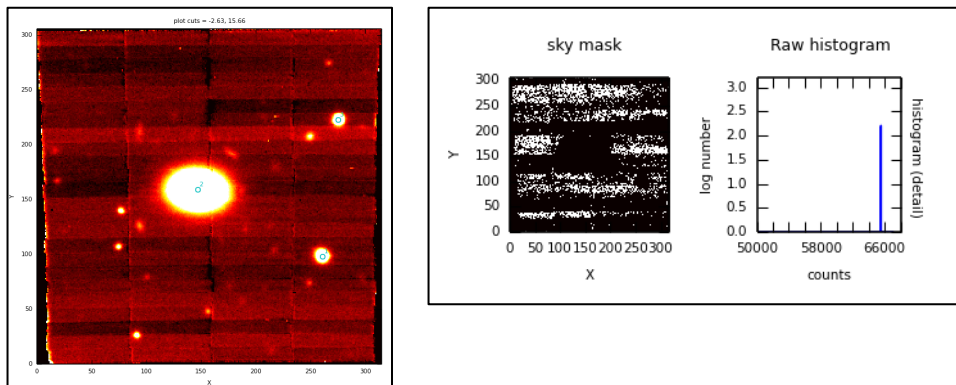


Figure 7. Datacubes without dedicated SKY pointing get their SKY determined on the exposure itself. The pipeline-generated sky mask (the pixels that are taken for the SKY continuum and lines) is then displayed on the QC report (bright pixels, right box), the dark pixels have been omitted. For comparison, the corresponding IMAGE_FOV file is displayed in the left box.




Figure 8. QC plot of one of the individual datacubes (for comparison to the plot for the COMBINED datacube, in case of alignment errors). It shows the same properties and parameters as the previous figure, except for: exptime (exposure time of the parent raw file); SKY_YN: Y if this datacube has used a dedicated SKY observation for SKY subtraction; Nsat = number of saturated pixels.

Known features and issues

1. Issues: Misalignment




Figure 9. Combined datacube with duplicated sources.

The alignment procedure (recipe `muse_exp_align`) aims at detecting cases of “telescope wobble”. This effect can happen with the instrument derotator, causing small unintended shifts between subsequent exposures that are not registered in the header. This effect is rare. The recipe attempts to detect a possible misalignment numerically, by extracting sources and matching them. The underlying assumptions are: the sources should be clearly detectable and should be the same in all exposures. They are a bit risky however.




Figure 10. Combined datacube (N=12) with alignment correction turned off. In this case there was real telescope wobble, with some input files badly combined.




Figure 11. The same datacube, with alignment enforced. This worked well in this case.

Therefore this step is generally not applied in cases of more than 4 input files which is typical for a mapping strategy with little or no overlap. It could also fail e.g. because the sources are faint or misinterpreted. If such a case is prominent enough to be discovered in the certification step (see Figure 9), the datacube gets re-processed without alignment. This is flagged by the value 1 in the alignment bit #10. There is a small chance that such issues might get overlooked.

The opposite case might happen as well (Figure 10 and Figure 11). With the alignment correction turned off by default, this N=12 combined datacube shows heavy misalignment, due to real instrument wobble being uncorrected. The same datacube, with the alignment correction enforced (Figure 11) is well aligned.

If a datacube looks suspicious:

- check alignment bit #10,
- compare the combined product QC plot to the individual QC plots,
- if confirmed: download the individual datacubes and use them for scientific analysis.

The procedure to get access to the individual datacubes is described under “[Special download channel](#)”.

Guide star lost.




Figure 12. Guide star lost.

Some exposures suffer from the guide star being lost (see example at left). This cannot be detected by the pipeline since *there is no information available about this issue in the raw file headers*. Often but not always this is marked as an OB comment. Often but not always this is picked by the QC review process. Please check carefully in the QC plots. If such a case is found, the combined datacube is normally not usable for scientific analysis. You may want to work with the individual datacubes instead, see ‘[Special download channel](#)’.

Residual rotation (instrumental error).




Figure 13. Rotational misalignment, combined datacube.




Figure 14. Rotational misalignment, single datacube.

Very rarely a sequence of exposures, designed to be rotated by 90 degrees each, seems to have an incremental rotational error between two exposures, not corresponding to the standard 90 degrees (or multiples thereof) but a few degrees more or less. This deviation is not recorded in the header. The alignment correction procedure of the pipeline, working on the nominal values, can only detect and correct for linear shifts. Then, no matter if the alignment correction is applied or not, the superposition will be wrong, as can be seen in the comparison between the combined image (Figure 13) and a single exposure (Figure 14). The difference to the linear displacement in Figure 9 is that the offsets are not the same across the image of the datacube, which makes them particularly difficult to detect. In case of doubts you may want to work with the individual datacubes instead, see [‘Special download channel’](#).

Background subtraction: general

For OBs with dedicated SKY measurement, check carefully the OB grade and comments whether the conditions were stable. The pipeline cannot take variable transparency or cloud conditions into account. Depending on circumstances, the applied SKY level might lead to over or under-subtraction.

There are two OB strategies for SKY subtraction: one with dedicated SKY pointings, and one without²⁰. Both have certain risks, and the user should be aware of those. For datacubes with dedicated SKY pointings, the sky level can be well extracted by the pipeline, but in order to safely apply it to the OBJECT (science) data, one must be sure that the conditions are stable, both in terms of photometry and moon visibility. This cannot be checked by the automatic processing scheme that we have used. Any issue here could cause negative (obviously wrong) background levels in the result datacube.

For datacubes without SKY pointings, the risk is that each of the individual exposures gets its own background level, or that the SKY is generally over-estimated (e.g. in crowded fields, see below), again resulting in negative flux values.


In case of doubts you may want to download the individual datacubes and check them, see [‘Special download channel’](#).

²⁰ We consider here the case of ‘hidden sky’ as a normal SKY pointing.

Background subtraction: ultrashort exposure times

Some OBs contain observations with very short exposure times, <10 sec, down to 0.02 sec. They have been designed to avoid saturation of bright sources in the field. First versions of the pipeline used to fail in such cases, so we decided to skip processing of those data in our standard scheme. However, with mature pipeline versions the situation has improved. As of P104 (October 2019), data with very short exposure times, <10 sec in WFM and <150 sec in NFM, are processed normally. Very short integration means that the sky lines are underexposed and pipeline tends to apply wrong differential wavelength correction based on them. Thus, we skip this step (parameter --skylines=NONE in recipe muse_scibasic) for those data.


Background subtraction: crowded fields



Some OBs have been designed by their PI such that they measure a crowded field, without taking a separate SKY field for proper SKY determination and subtraction (Figure 15, top). In that situation (which is not known *a priori*) the reduction strategy is that the pipeline derives a sky fit from the pixels corresponding to the 20% lowest counts (parameter `skymodel_fraction`). In those cases without extra SKY, the sky mask as calculated by the pipeline is displayed in the QC report. The pipeline strategy might go wrong in such a case since it likely includes the lower parts of the stellar PSFs, over-estimating the background and then over-subtracting the data (maximum value in the example histogram at -40, broken line). The QC procedure discovers those cases *a posteriori* and marks them by setting flag #8 to 1 (bottom).

The observational strategy in this example likely did not aim at precision photometry but instead at RVs (where the negative fluxes don't matter). Nevertheless it should be possible to obtain a reasonable photometry, by applying aperture photometry on the sources, whereby the negative offsets for the source and for the local background would compensate each other.

Figure 15. Crowded field, back-ground over-subtraction. From top to bottom (extracted from the QC plot): the field of view, the sky mask (bright marks the input used for the sky), the product histogram showing the negative background (maximum marked by the broken line), and the labelling indicating the issue (histo_mode, score_bit, SKY_YN=N).



Another example demonstrates that a datacube with extra SKY pointing receives a decent sky subtraction, and a reliable background level, even for a crowded field (Figure 16).

Figure 16. Crowded field, good background subtraction.

ABMAGLIM. This parameter needs to be interpreted with some care. In the cases of over-subtracted background (in particular crowded fields without extra SKY) it is underestimated. Even with proper background subtraction, this parameter would measure, in crowded fields, the background fluctuations induced by faint stars.

Quality issues for the first Science Verification period. For the first SV period in June 2014 we have encountered quality issues with the calibrations. In particular for dates 2014 June 23, 24, 25, and 27, the illumination and tracing calibrations did not match well the prevailing conditions of the science data, due to particularly low ambient temperature. IFU #6 was affected, but also the other IFUs. We reduced the science data with master calibrations from dates 2014 June 22 and from June 28, respectively, which are intrinsically of good quality but do not match perfectly well the conditions on the science dates. We decided to include the reduced data nevertheless in the release but users should be aware of their limited quality.

Wiggles in AO-E data. For AO-E data (extended mode), the current pipeline version does not correct properly for bumps and wiggles in the instrument response function. These are caused by the transmission of the Na filter that is used to block the contamination from the AO laser star. The wiggles are propagated to the science spectra. They are particularly evident in the blue part of the spectrum (see Figure 17).




Figure 17. Wiggles in the spectra of for instrumental setup AO-E.

Raman scattered laser lines. The MUSE observations using the AO system are affected by the Raman-scattered light from the lasers. Its contribution is seen mainly as emission lines at 6485 Å and 6827 Å, which fluxes vary slowly across the field of view by about 5%. Currently, the MUSE IDPs are not corrected for the Raman lines.

Quality issue with data obtained between February and April 2019. Due to an issue with the telescope dome safety camera, MUSE observations taken between February 1st and April 18th, 2019, suffered from light contamination. The contamination is visible as an excess of continuum emission between 800 and 900 nm, with the peak around 860 nm (Figure 18, Figure 19). The rest of the spectral range is not affected.




Figure 18. Example of the red part of MUSE spectra of the same program and the same pointing, where data were taken before (blue) and during (black) the affected period. The contamination is visible as a bump between 800 and 900 nm peaking at 860 nm.




Figure 19. Examples of the contamination residuals in the red part of spectra of different sky regions (darkest regions of the field of view) of a single data product. Note the variable contamination residuals between 800 and 900 nm.

The analysis of data products from that period shows that contamination might vary from exposure to exposure. This is propagated to the combined, final data product. In some cases, when the sky background is estimated on the same science frame, the pipeline successfully removes contamination from the final data cube. However, when the sky background is created from a dedicated SKY pointing, there are often residuals of contamination in the final products, presumably due to the variability of the sky.

Moreover, the residuals of contamination in MUSE data products vary across the field of view (Figure 19). This is probably due to variability in the light source itself and to possible reflections within the telescope dome. It also may be due to strategies applied in automatic pipeline processing.

Users are urged to use the MUSE IDPs from that period of time with caution. Any residuals of the contamination should be taken into account for the analysis of the scientific signal.

Quality issue with data obtained between June 15 and June 19, 2022.




Figure 20: Example of the vignettted field-of-view due to camera beam splitter in the optical path.

The MUSE on-sky: science, standard star and twilight flat observations, acquired in that period of time, show vignetting (see example). It was caused by the GALACSI commissioning camera beam splitter inserted in the optical path. In the WFM, this element defines a square 82"x82" within which the light that is passed to MUSE is decreased by approximately 50% at 500 nm to 75% at 900 nm. The light modification may vary across the field-of-view. The vignettted master calibrations were used to process vignettted science exposure.

Users are urged to consider the MUSE IDPs from that period of time with caution.

Mosaic with small overlap between different pointings




Figure 21: Example of the mosaic consisting of two different pointings. Note the small overlap area between them, with very few sources. This affected the alignment found by the automatic procedure and resulted in suspicious looking datacube of the mosaic.

Some OBs are designed by their PI as a mosaic of two (or more) different pointings with small overlap between them. For each pointing, there are several exposures with many sources in the field-of-view, which allows the alignment procedure to find a good alignment for combination. However, the overlap area(s) is very small and/or lacking sources that can be used to find the proper alignment. In such cases, although the individual pointings are combined with good alignment, the alignment of the mosaic cannot be guaranteed.

If a datacube of the mosaic looks suspicious, download the individual datacubes and use them for scientific analysis. The procedure to get access to the individual datacubes is described under [“Special download channel”](#).

2. Features




Figure 22. Spurious sources claimed by the pipeline, with no real meaning.

Parameter Nsources. The pipeline extracts sources on the IMAGE_FOV file, in order to measure their FWHM and derive QC parameters for image quality. They are also plotted on the QC plots. The number given is very conservative. It applies to point sources only. It is sometimes spurious: for extended objects, many pseudo sources are claimed that seem to have no real meaning (Figure 22).

Residual shifts of wavelength scale. If QC flag #9 is set, the absolute value of the maximum residual shift derived from sky lines and applied in the muse_scibasic step is larger than 0.2 \AA (if processed before October 2019), or 0.5 \AA as of P104 (October 2019). This is higher than usual, and the resulting wavelength scale should be checked carefully.

Flux calibration. MUSE datacubes are not necessarily taken under photometrically stable conditions. User constraints might be relaxed (THIN or even THICK clouds permitting), or user-defined constraints might have failed (OB grades C or D), with the datacube nevertheless processed. Combined datacubes might therefore have flux-related properties and noise varying across the datacube. The combination recipe does take varying exposure times into account (although these are rare) but cannot account for varying flux levels.

For the flux calibration, the closest-in-time standard star has been taken, usually from the same night. However there is no guarantee that the photometry resulting from the standard star is applicable to the MUSE data products, and the flux calibration should be considered to be relative. In non-photometric nights, the flux calibration can go wrong in both directions: the flux scale of the datacubes could be over-estimated, under-estimated or correct, depending on the relative extinction differences between the science and the standard star.




Figure 23. Histogram for a datacube with saturated data.

Saturation. Check carefully the saturation flag #6. If set to 1 in a combined datacube, at least one of the input files has an issue with saturated pixels. In a combined datacube it might be difficult to tell which spaxels got affected. The QC plots of the individual datacubes might give further information about the level of saturation. If saturated pixels occurred, be very cautious with the analysis. Pixels might even show negative counts in extreme saturation. You may want to check the individual datacubes from the [“Special download channel”](#).




Figure 24. Datacube contaminated by satellite trail.

Satellite trails. Occasionally one out of N input datasets has a satellite trail which propagates into the combined datacube. These datacubes are left untouched. If a user finds that artefact disturbing, there is a possibility to download the individual datacubes and analyse them individually. See ‘[Special download channel](#)’ for more.




Figure 25. Combined datacubes with rotational offsets might have cross-like artefacts.

Crosstalk. In case of saturation, some CCD columns might suffer from ‘bleeding’. The typical rotation pattern by 90 degrees affects different CCD columns and causes cross-like features in the combined datacube. The individual datacubes have these artefacts minimized, in one direction only. This is another example when a user might want to inspect the individual datacubes. See ‘[Special download channel](#)’ for more.

Ring-like features. Rarely there are faint ring-like features visible, apparently centered. They are related to reflections of the M2 structure on twilight flats and cannot be removed by the pipeline.

PI-specified constraints. It might be useful to know the user-specified constraints for an observation, for Service Mode OBs. In particular, the transparency constraint is stored in the key “HIERARCH ESO OBS AMBI TRANS”.

Tips and tricks

Post-pipeline removal of sky lines. The MUSE data cubes are sky subtracted but may still include residuals of sky emission lines rewards of $\lambda > 7000\text{\AA}$. These residuals can be further reduced using an `ESOReflex` interactive workflow, to be applied directly to the MUSE science data products. It is called `muse_zap`. It uses the Zurich Atmosphere Purge code²¹ and is available, together with the MUSE pipeline and other workflows, from the ESO Instrument Pipelines webpage²². There you also find a tutorial, plus detailed instructions on the installation procedure and usage of the workflow.

Analysis software package. For their data analysis, users may be interested in MPDAF, the *MUSE Python Data Analysis Framework*. It is available under <http://mpdaf.readthedocs.io/en/latest/index.html>. This is an open-source (BSD licensed) Python package, developed and maintained by the *Centre de Recherche Astrophysique de Lyon*²³.

Special download channel.

Purpose. Our QC process involves an automatic scoring and a visual inspection. We believe that we find most of the cases where an alignment correction failed, or when an input dataset contained a bad file (guide star lost). If however a case of bad input file or of bad combination got undetected, you can visually inspect the `png` files and the delivered combined datacube.

²¹ Soto et al. 2016, MNRAS, 458, 3210 [<http://cdsads.u-strasbg.fr/abs/2016MNRAS.458.3210S>]

²² <http://www.eso.org/sci/software/pipelines/>

²³ <https://cral.univ-lyon1.fr/?lang=en>

It might be useful in some cases to then have access to the SINGLE datacubes. These could be used for an independent check or even for a scientific analysis. While it is true that the combined datacube offers best possible contrast, an analysis of the single datacubes might be a valid alternative in an environment which is limited in terms of the hardware necessary to fully reduce these data. We therefore offer a special channel for downloading them.

Procedure. Check the header of the combined datacube. It will contain the names of all the single datacubes, in two syntaxes: one is the 'original filename' DCORG<n> (the one used in the data reduction system), the other is the 'archive filename' DCARC<n> (the one used in the archive). <n> is the file index.

You will need the archive filename under DCARC<n>, e.g. DCARC1. Now go to the ESO archive interface for direct file retrieval²⁴ and enter the name under DCARC1 (without the extension .fits). Add the other DCARC<n> filenames, and then push 'Request Data Files', provided that you have the permission to access these files²⁵. For convenience you may then want to rename the downloads from their DCARC names to the DCORG names. The individual datacubes will have a size of 3 up to 5 GB each.

Working with the single datacubes.

With the individual datacubes downloaded in case of an issue with the combined datacubes, one may continue analysis and/or further combination using <https://mpdaf.readthedocs.io/> or <http://ifs.wikidot.com/>. Please keep in mind that combining the individual datacubes is in principle less accurate than using the combined datacube which has been resampled only once. As mentioned above, this strategy is only the secondary choice, if there is an issue with the combined datacube.

Working with pipeline log files.

The associated pipeline log file is delivered with every MUSE datacube. It contains the log files from all pipeline steps and reduction jobs that contributed to the datacube. In particular, the log file can be parsed to check

- what corrections have been applied,
- which pipeline parameters have been used,
- correction values (like barycentric RV correction).

The first part has an overview of the processing job cascade, with the involved recipes. The second part lists all pipeline SOF (set of input frames) and pipeline logs, sorted in chronological order.

To find the complete log of a particular recipe 'rblog' in the complete pipeline log file 'logfile' of a particular datacube, type

```
grep "<rblog>" <logfile>
```

Each recipe log starts with some information about the pipeline version, the used non-default parameters, and the entire recipe call.

²⁴ <http://archive.eso.org/cms/eso-data/eso-data-direct-retrieval.html>

²⁵ These are the same as for the combined datacube.

Data Format

Files Types

The primary MUSE product is the 3D datacube which could come in 4 different flavours:

ORIGFILE names starting with	Product category HIERARCH.ESO.PRO.CATG	Format	How many input files?	Description
MU_SCBC	DATA_CUBE_COMBINED	3D image	N>1	combined datacube from OBJECT observations
MU_SCBS	DATA_CUBE_FINAL		N=1	datacube from a single OBJECT observation
MU_SCBY	DATA_CUBE_COMB_SKY		N>1	combined datacube from SKY observations (<i>shallow</i> datacube)
MU_SCBX	DATA_CUBE_FINAL_SKY		N=1	datacube from a single SKY observation (<i>shallow</i> datacube)

Each product has exactly one ancillary FITS files delivered with it:

ORIGFILE names starting with	Product category HIERARCH.ESO.PRO.CAT G	Format	How many input files?	Description
MU_SIMC	IMAGE_FOV_COMBINED	2D image	N>1	collapsed white-light image of combined FOV, from OBJECT observations
MU_SIMS	IMAGE_FOV		N=1	collapsed white-light image of FOV, single observation
MU_SIMY	IMAGE_FOV_SKY_COMB		N>1	collapsed white-light image of combined FOV, from SKY observations
MU_SIM	IMAGE_FOV_SKY		N=1	collapsed white-light image of FOV, from single SKY observations

Furthermore the following non-FITS files are delivered with each datacube:

ORIGFILE names starting with	ASSOCn (listed in the header of the main product)	Format	How many?	Description
If the product is a COMBINED datacube made of N input files:				
r.MUSE...tpl.png	ANCILLARY.PREVIEW	png file	1	See Figure 6.
plus: r.MUSE...pst.png	ANCILLARY.PREVIEW	png file	N	See Figure 8.
r.MUSE...tpl.log	ANCILLARY.README	text file	1	all recipe processing logs for the combined datacube
If the product is a SINGLE datacube:				
r.MUSE...pst.png	ANCILLARY.PREVIEW	png file	1	See Figure 8.
r.MUSE...pst.log	ANCILLARY.README	text file	1	all recipe processing logs for this datacube

The following naming convention applies to the ORIGFILE product: e.g. the name MU_SCBC_1117772_2015-04-12T00:56:47.087_WFM-NOAO-E_OBJ.fits

has the components:

ORIGFILE component:	MU	SCBC	1117772	2015-04-12T00:56:47.087	WFM-NOAO-E_OBJ.fits
refers to ...	MUSE	product type (SCB stands for science)	OB ID	timestamp of first raw file	setup string: wide-field mode, no AO, extended

ORIGFILE component:	MU	SCBC	1117772	2015-04-12T00:56:47.087	WFM-NOAO-E_OBJ.fits
		cube; C – combined, S – single, Y – combined SKY, X – single SKY)			wavelength range; DPR.TYPE=OBJECT (could also be SKY)

The ancillary files have the following ORIGFILE names:

Table 9. Naming conventions of ANCILLARY files

type	example	rule
ANCILLARY.README	r.MUSE.2015-04-12T00:56:47.087_tpl.log	Technical filename of the main fits file, with extension log instead of fits
EITHER (if COMBINED datacube):		
ANCILLARY.PREVIEW	r.MUSE.2015-04-12T00:56:47.087_tpl.png	same name, with extension png instead of log
ANCILLARY.PREVIEW (N)	r.MUSE.2015-04-12T00:56:47.087_pst.png, ...	names of all individual exposures
OR (if single datacube):		
ANCILLARY.PREVIEW	r.MUSE.2015-04-12T00:56:47.087_pst.png	name of single exposure

The user may want to read the ORIGFILE header key and rename the archive-delivered FITS files accordingly.

File structure

The MUSE datacube product has two 3D image extensions:

- 3D datacube with 2 spatial dimensions and 1 wavelength axis, with flux-calibrated spatial pixels; the EXTNAME key is 'DATA'.
- 3D datacube with the variance, EXTNAME is 'STAT'.

File size

The size of a *single* datacube is 3-5 GB. The same size applies to a datacube combined from datasets with small offsets. The size grows in proportion to the number of non-overlapping pixels. The larger values apply to datacubes with orientations inclined with respect to the RA/DEC grid. The size of a datacube made of 4 non-contiguous input files is about 12-20 GB.

Processing a combined datacube from 4 input files requires a memory of about 100 GB. Combining a datacube of 12 input files takes about 3 times as much, 300 GB.

Acknowledgment text

According to the ESO data access policy, all users of ESO data are required to acknowledge the source of the data with an appropriate citation in their publications. Since processed data downloaded from the ESO Archive are assigned a Digital Object Identifier (DOI), the following statement must be included in any publications making use of them:

Based on data obtained from the ESO Science Archive Facility with DOI(s) :
<https://doi.eso.org/10.18727/archive/41>.

All users are kindly reminded to notify Mrs. Grothkopf (esodata at eso.org) upon acceptance or publication of a paper based on ESO data, including bibliographic references (title, authors, journal, volume, year, page numbers) and the observing programme ID(s) of the data used in the paper.