

**On-board software maintenance operations analysis for modern LEO satellites:  
EarthCARE mission through its commissioning  
Marco Casanova-Álvarez<sup>a\*</sup>, Isabel Rojo<sup>b</sup>, Gabriela Ansteeg<sup>c</sup>, Javier Larrea<sup>d</sup>, Andreas Jung<sup>e</sup>**

<sup>a</sup> *EarthCARE DHS Spacecraft Operations Engineer, Serco Services GmbH c/o ESA-ESOC, European Space Operations Centre, Germany, marco.casanova.alvarez@ext.esa.int*

<sup>b</sup> *EarthCARE, SMOS, FORUM Missions Operations Unit Head, Earth Observation Missions Operations Division, European Space Operations Centre, Germany, isabel.rojo@esa.int*

<sup>c</sup> *EarthCARE Spacecraft Operations Manager, Earth Observation Missions Operations Division, European Space Operations Centre, Germany, gabriela.ansteeg@esa.int*

<sup>d</sup> *EarthCARE ATLID Spacecraft Operations Engineer, Telespazio Germany GmbH c/o ESA-ESOC, European Space Operations Centre, Germany, javier.larrea@ext.esa.int*

<sup>e</sup> *Senior System Software Engineer, Flight Software Section, European Space Research and Technology Centre, Netherlands, andreas.jung@esa.int*

\* Corresponding Author

**Abstract**

This paper presents the results of the on-board software maintenance (OBSM) operations performed during the commissioning phase of EarthCARE, a joint mission between ESA and JAXA that aims to study the interactions between clouds, aerosols and radiation. On-board software maintenance operations are generally constrained by the data uplink rates from ground to the spacecraft, on-board command queue limitations, hardware specific constraints, station visibility constraints, and/or radiation hardness requirements. In general, they constitute very critical operations.

The study presented in this paper focuses on the different OBSM patching operation strategies adopted for large and small software updates during the EarthCARE commissioning phase. These strategies have been developed as the result of different trade-offs to minimize effects due to high radiation areas, commanding link losses, memory corruptions and operator errors, all of which provided lessons learned that have contributed to developing a more robust OBSM overall strategy.

The paper also discusses the potential applications and benefits for future LEO missions, especially those with large constellations and/or agile software development life cycles. Finally, recommendations and best practices for planning, executing and verifying OBSM operations based on the experience gained from EarthCARE are presented.

**Keywords:** operations, on-board software management, LEO satellites, OBC patching

**Acronyms/Abbreviations**

ANX	= Ascending Node Crossing
AOS	= Acquisition of Signal
ASW	= Application Software
ATLID	= Atmospheric Lidar
CSW	= Central Software
CPDU	= Command Pulse Distribution Unit
EEPROM	= Electrically Erasable Programmable Read-Only Memory
ESA	= European Space Agency
EC	= EarthCARE
ESOC	= European Space Operations Centre
FOS	= Flight Operations Segment
KIR	= Kiruna ground station
HDSW	= Hardware Dependent Software
HKTM	= House-keeping Telemetry
HPC	= High Priority Command

JAXA	= Japan Aerospace Exploration Agency
LEO	= Low Earth Orbit
LEOP	= Launch and Early Orbit Phase
LTDN	= Local equator crossing Time on Descending Node
MCS	= Mission Control System
MTL	= Mission Timeline
OBC	= On-board Computer
OBSM	= On-board Software Maintenance
PDGS	= Payload Data Ground Segment
PUS	= Packet Utilization Standard
RT	= Real-time
SAA	= South Atlantic Anomaly
SAU	= Smallest Addressable Unit
SC	= Spacecraft
SGM	= Safe-guard Memory
SSO	= Sun-synchronous Orbit
SW	= Software
TR	= Troll Ground Station
TT	= Time-tagged

## 1. Introduction

The EarthCARE mission was launched on 28 May 2024 on a Falcon 9 launch vehicle from Vandenberg, United States, and is operated by ESA's European Space Operations Centre (ESOC) in Darmstadt, Germany. EarthCARE is the largest and most complex ESA's Earth Explorer mission to date. It employs high-performance lidar and radar technologies, which have never before been flown together in space, with the objective of providing unprecedented datasets. These datasets will also combine data from the on-board radiometer and imager instruments to study the relationship between clouds, aerosols, and radiation. This type of instrumentation allows to have vertical profiles of several meteorological and climatic variables whereas before, only top of the atmosphere readings were available.

EarthCARE (EC) follows a sun-synchronous near circular frozen orbit (SSO) with a mean altitude of 393.14 km and an inclination of 97.05°. The orbit has a period of 92.5 minutes and a repeat cycle of 25 days, having its Local equator crossing Time on the Descending Node (LTDN) at 14:00 hours. The orbit altitude is set as low as possible to optimise the performance of the active instruments while keeping propellant requirements for orbit maintenance within acceptable limits. The drag environment around the Earth has a high level of variation due to unpredictable solar activity, and at 400 km altitude the quantity of drag is severe. A direct impact of the reference altitude for operations is the reduced ground station visibility, which provides at most ~8 minutes of tracking above 5 degrees of elevation (limit set for commandability).

While Science data and house-keeping are downlinked in X-band frequency at a rate of ~150 Mbit/sec, command data and also house-keeping are transmitted via an S-band transponder. The downlink is performed at nominal rate of 2 Mbit/sec (128 kbit/s in low bit rate mode, which allows ranging) and the uplink at 64 kbit/sec.

EarthCARE's on-board computer (OBC) is cold redundant, with one active processor and one switched off (which can power on in parallel and is handled via an interlink processor communication link to the active processor). Each of the processors have an Electrically Erasable Programmable Read-Only Memory (EEPROM) divided into two banks, allowing to have redundant images. This plays a key role specially during OBSM operations as the opposite bank to the one being patched can be preselected in case of a boot process. In terms of size, EC's EEPROMs are 4 MiB (2 MiB per bank), which is larger than those used in previous Earth Explorer missions such as Swarm, CryoSat, or Aeolus.

## 2. EarthCARE Operations overview

Due to EC's SSO, the mission relies on S-band ground stations positioned near the poles. In the Northern Hemisphere, it utilizes ESA's Kiruna Station (KIR) in Sweden and KSAT's Svalbard Station (SG) in Norway. In the Southern Hemisphere, it uses KSAT's Troll Station (TR) in Antarctica. This approach allows to have visibility of the satellite twice per orbital period, which has been necessary during critical operations such as the Launch and Early Orbit Phase. After commissioning, the required commanding activities are considerably reduced, allowing to have only 2 S-band passes per day (with both downlink and uplink capabilities) and 10 X-band passes with downlink capability to both downlink house-keeping telemetry (HKTm) and scientific data from the four instruments.

The satellite's operational concept is highly oriented towards autonomy from ground. Each week, mission planning is conducted to schedule activities requested by the Payload Data Ground Segment (PDGS) for each instrument, as well as tasks related to the Flight Operations Segment (FOS), such as status reporting, S-band pass commanding, and other essential operations. The plan is then reviewed to ensure there are no conflicts between the different activities, nor violations of operational constraints. The outcome of it are commands to be uplinked to the spacecraft with a concrete time or orbital position at which they will be executed. This time or angle execution target is defined in the Packet Utilization Standard (PUS) Services 11 (Time-based scheduling, already present in PUS-A) for the mission timeline (MTL) and a precursor of PUS-C Service 22 for position-based scheduling [3].

In each of the passes, a special buffer of the safe-guard memory (SGM) is downlinked to see if any event of high criticality (Service 5) has been generated since the previous contact. If this is the case, the buffer is also cleared to free space for future anomalies.

## 3. Constraints for OBSM operations

OBSM operations, especially for the OBC, constitute one of the most critical activities that can be performed in a space system, not only by its importance but also by its challenges or risks. There is a large experience on OBSM at ESOC, ranging from interplanetary missions [4], [5] to LEO regimes [6].

In the case of EarthCARE, the OBC EEPROMs that are used to store the central software (CSW) of the spacecraft are equipped with a write protection mechanism. This allows the application to lock the memory content by a Hardware Dependent Software (HDSW) service. This feature does not only prevent operator errors overwriting memory addresses but it also reduces the risk of memory corruptions due to radiation events. During OBSM operations, before writing any content, it is required to disable the write protection. As a consequence, there is an increased risk of memory corruption during the operation duration.

While the simplest way of patching an on-board memory would be during visibility in real time, previous Earth Explorer Missions at ESOC have performed OBSM operations normally using one pass to uplink to the MTL the patch commands with an execution time lying inside the next pass in order to execute in visibility and being able to disable and enable the write protection before and after the execution of the patch commands. This allows to decouple the data transfer from the operation execution, making it safer as the write protection would only be disabled after ensuring that the full batch of patch commands is on-board with no data transfer loss experienced. However, the fact that the ground stations used during the routine phase are located at the poles makes it difficult to minimize memory radiation doses during contacts.

The on-board TC buffer is split into a given number of slots for time- or orbit-position tagged commands depending on the size of the commands. The overall buffer is split into three: small, medium and large size TC buffers. Depending on the patch strategy (e.g. to uplink full image, individual non-consecutive differences), this may constitute a constraint of the number of commands per pass on which one can divide a patch operation.

Figure 1 shows the visibility of EarthCARE from the ground stations with respect to the ascending node crossings (ANX) as well as the South Atlantic Anomaly (SAA) crossings. Due to the sun-synchronous nature of the orbit, this pattern of visibilities by time of the day is constant along the year. This figure has been produced with the objective to spot the most beneficial regions and ground contacts to minimize the effect of radiation during OBSM operations. Due to their locations, EarthCARE's ground stations are in areas of high radiation dose. However, the highest radiation

levels are experienced in the SAA anomaly crossings, happening between midnight and 8:30 am before crossing the ascending node of the orbit or between 11:00 am and 20:00 pm between KIR and TR visibilities. Times are given in UTC.

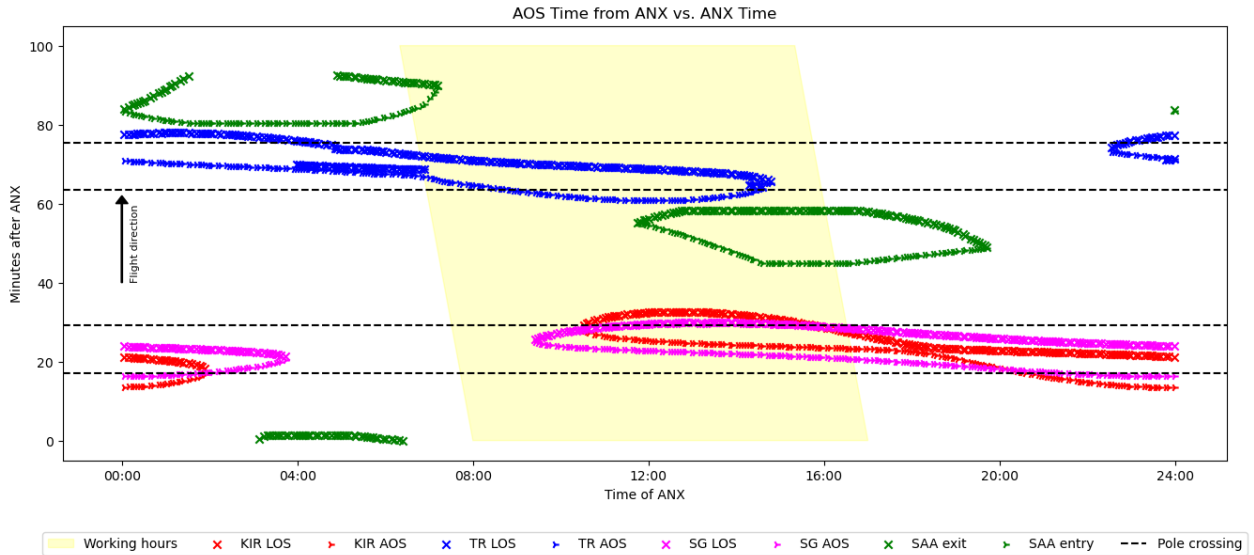


Figure 1 EarthCARE ground contacts visibility distribution.

Due to the importance and criticality of an OBSM operation, visibility during the patch is preferable but not required. In the case of EarthCARE, it is possible to issue specific high priority commands (HPCs) or hardware commands from the active processor as software commands sent by ground using Service 2 instead of the regular HPC sent from Ground which are forwarded by the TC decoder directly to the command pulse distribution unit (CPDU) without involvement of software (SW). This allows to send some of these types of commands outside of visibility which is important in the scope of an OBSM operation as the enabling and disabling of the EEPROM write protection commands are of this kind. Which commands can be executed in this mode in EarthCARE will be discussed in the following sections.

#### 4. Design, planning and execution of OBSM operations in EarthCARE

Since launch, there have been already three different OBSM operations conducted in EarthCARE. Each of them differs from the others in size, criticality and units involved, and different design criteria have been used to define each of the operations. The aim of this section is to explain the trade-offs taken based on the above mentioned constraints to share the advantages found in different patching strategies.

In June 2024, a patch was performed in the atmospheric lidar (ATLID) instrument for performance optimization of the laser. In July 2024, a small patch of the CSW of the order of bytes was required to improve thermal aspects of ATLID in case of a software reboot. Finally, in April 2025, a large patch size was performed for the CSW to improve performance aspects of the propulsion system.

##### 4.1 ATLID patch operation

The ATLID Application Software is distributed across two EEPROM memories due to its size. The OBSM operation was the update of the ASW, which on a first approach would have required the full patch of both memories.

To optimize this operation, the differences between the existing ASW EEPROMs and the updated version were analysed. However, for this particular ASW delivery, the differences between the previous and the new compiled versions were substantial. The number of commands required to uplink individually each non-consecutive change between uplink and base software images was higher than the number of commands of biggest size required to simply uplink the whole new image. Given the granularity of the changes, which affect individual smallest addressable units

(SAUs), it was determined that analysing them globally would be more efficient than using a patch approach based on differences between the two compiled software versions at the individual change level. This latter approach is what the OBSM tool of ESA’s mission control system (MCS) provides by default.

The final approach for reducing the number of commands needed to fully patch the ATLID EEPROM with the new ASW was to identify the largest memory areas that remained unchanged and eliminate commands containing data related to those areas. The following figure illustrates the process followed:

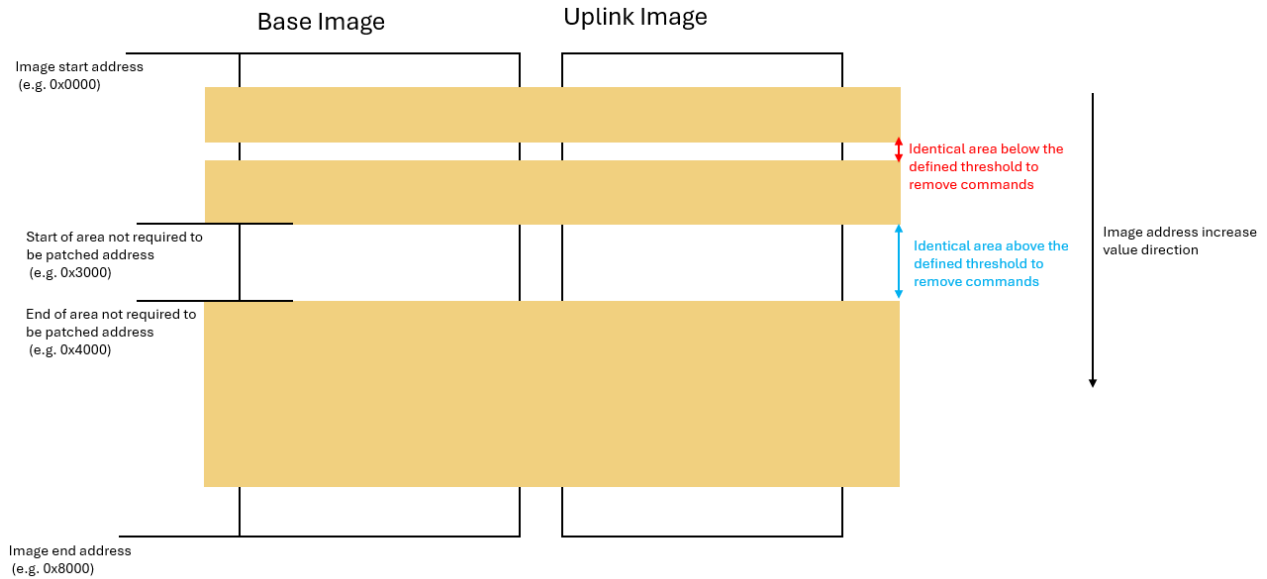


Figure 2 Diagram of memory addresses with commonalities or differences between base and uplink images depending on the size of the identical areas. Yellow areas represent areas with differences of content.

First, a stack with all the commands of long size required to completely patch the EEPROM with the new image was generated. As a second step, the commands which were only comprising memory addresses contained in areas of identical values higher than a certain threshold, have been removed to reduce the number of commands required. In this case, the threshold was set to 9 KiB.

This reduction of commands has comprised 11 areas of the memory and a total of 0.5 MiB worth of data. It allows to efficiently select only the areas where a patch is required from a global perspective as treating individually each non-sequential change would have led to 4 times the number of required commands. Moreover, this optimised approach was translated into a reduction of passes required for the operation which also led to a cost and effort reduction. With the final number of required commands, they were split dividing them by the maximum number of patch commands to be used per pass. This number is based on the size of the long size TC buffer of the spacecraft and a certain margin to allow other operations to use long size commands.

Contrary to the OBC EEPROMs, the write protection implementation of ATLID is a software command rather than an HPC. This allows to add these commands to the time-based queue of the spacecraft which then does not constraint at which point of the orbit the operation is taken place. As seen in Figure 1, by the polar latitudes of our ground stations, based on radiation impact risks, it is safer to perform this operation towards equatorial latitudes outside of the SAA crossings. Moreover, this provides more time within the pass in case of any contingency (e.g. ground station issues). Figure 3 shows how this patch has been performed outside of visibility conditions:

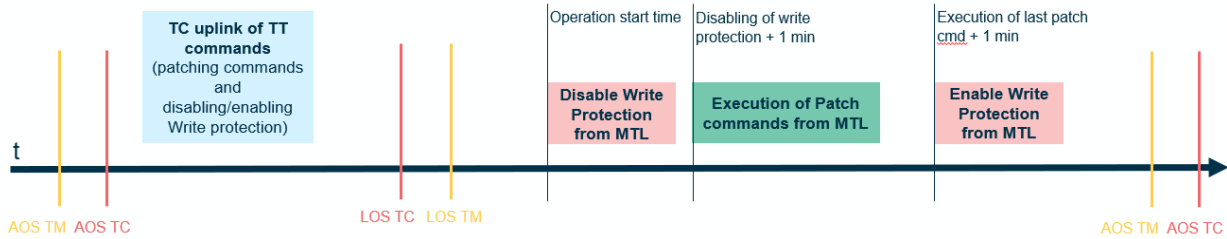


Figure 3 Patch concept followed for ATLID EEPROM.

During the contact, after ensuring that the spacecraft is in nominal state and that commanding capabilities are established after acquisition of signal (AOS), the patch commands were sent to the on-board commanding queues. Once the correct transmission of all the commands of each pass's chunk of patch commands was verified, the disabling and enabling of the write protection commands were also transmitted.

#### 4.2 CSW small patch operation

This patch consisted only of 20 bytes of differences split into 3 chunks of different start addresses. As the non-consecutive differences found when comparing the compiled versions of the base and uplink images were very reduced, only small commands containing the different SAUs were considered. It was also decided to simplify the operation by performing it in real-time (RT) during a single pass as the time needed to do so was as maximum half of the minimum pass commandability time. Figure 4 shows this concept, based solely on RT in contrast to previous fully time-tagged (TT) approach:



Figure 4 Patch concept followed for CSW small binary update.

#### 4.3 CSW large patch operation

In the frame of a propulsion system optimisation, a new CSW version was created. This last compiled CSW image received differed considerably from the previous one. The number of commands required to uplink individually each of the non-consecutive changes is more than double than the number of biggest size commands required to patch the whole OBC EEPROM content. As the individual changes mostly fit into a small size command (as they are only a few SAUs), the maximum number of commands which could be uplinked in each pass would be limited by the MCS release capacity and the TC decoder reception limit. After comparing this limit and the previously introduced limit of on-board long-size TC buffer, it was decided that similarly to ATLID's EEPROM patch, commands using their biggest size without considering individually differences among images was the best option. Then, as for the ATLID patch case, and as highlighted in Figure 2, the commands comprising identical memory areas among previous and new image bigger than 11 KiB were removed from the baseline. In this specific case, the approach led to a reduction of 60 % of the commands with respect to a difference-by-difference approach.

For the uplink strategy, the enabling and disabling of the write protection of the OBC's EEPROMs are HPCs due to the criticality of the units. Nevertheless, the re-enabling of the write protection can be performed encapsulating the

hardware command into a SW command via Service 2 (Device command distribution service) to be able to time-tag it. Due to this spacecraft design on which the disabling of the write protection was decided to be limited to visibility scenarios only, the operation is performed in partial visibility conditions, being a mixture of the two previous cases:

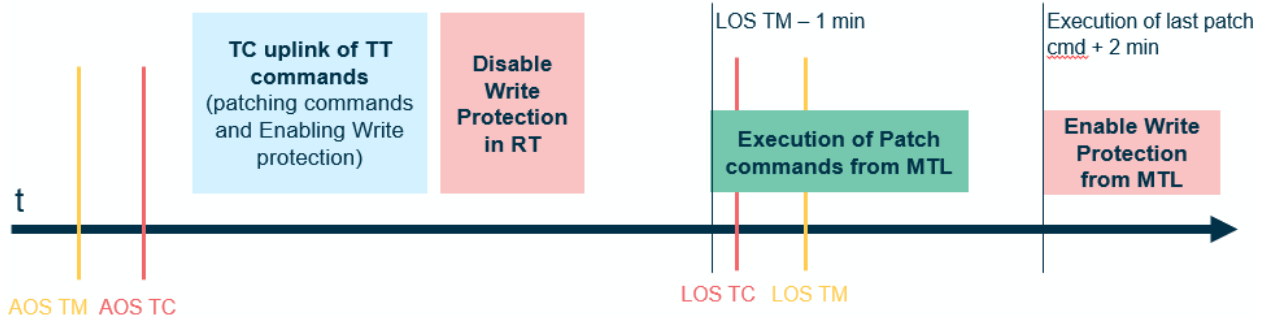


Figure 5 Patch concept followed for CSW large binary update.

This concept allows to use all the commanding time of the pass for uplink purposes (allowing to react to unforeseen events in the ground or spacecraft segments) and have partial visibility of the operation. Moreover, being executed right after the pass, for many cases, especially the passes with ground stations in the north, means to be outside of high radiation areas. The philosophy of uplinking first the commands and once confirming the success of the transfer removing the write protection also makes the operation more robust as in the case of contingencies, even if on-board, the commands will be simply rejected with no side effect.

## 5. Lessons learned

From the previous patching experiences, the following lessons learned have been identified. They have been divided into the two main sub-processes of an OBSM operation. The first sub-process is the command generation, which comprises the design of the commands to be sent to the spacecraft to perform the patch. The second sub-process is the uplink strategy, on which the visibility conditions and the way the commands are uplinked are studied.

### 5.1 Command stack generation

As explained in previous sections, once a new image is received, the differences with respect to the one on-board the spacecraft should be analysed to minimise as much as possible the bytes to be written and the amount of commands as the first one has an impact in the EEPROM lifetime and the second, in the uplink strategy. From the previous experiences, the following lessons learned were compiled:

- From the operational point of view, a deterministic build process to generate image binaries performing isomorphic mapping could reduce considerably the number of SAUs modified between two images. The ultimate impact of this would be a reduction in the complexity and scope than a patch operation requires. Specially for satellite constellations where a software upgrade would need to be applied to a large number of banks or missions with limited ground contacts these practises are advisable.
- Unless the amount of non-subsequent differences found between base and uplink image are at least smaller than the commands required for a full patch, differences should be studied and understood globally and not at microscopic level.
- MCS design should be flexible in terms of stack management for large command size and number operations such as this one.

### 5.2 Uplink strategy

An uplink strategy is required to divide the resulting command stack into ground contacts and to assess the degree of visibility needed for the operation as well as how it will be done (e.g. TT vs RT). The lessons learned on this part are the following:

- Due to the criticality of this operation, visibility conditions are an asset. However, they should not represent a constraint as for SSO satellites this would go along with the execution of OBSM operations in higher radiation dose areas and depending on the orbit's altitude, the ground contact times might also play a key role.
- Unless the number of commands is very small and the uplink time is at least less than half of a ground contact, it is advised to decouple transmission of the commands from its execution to be able to verify first the reception and then concentrate on the execution.
- The uplink strategy should also consider possibilities of losing passes completely or partially (i.e. only partially uplinked stacks) for rapid recovery and alternative completion of the uplink.

### 5.3 Workflow for optimising software patching operations

Based on the OBSM operations performed during the EarthCARE mission's commissioning phase and the collected lessons learned, Figure 6 provides the recommended workflow to optimise as much as possible a patching operation based on the images differences and satellite specific constraints.

## 6. Conclusions

This paper presents the different constraints and solutions adopted by the EarthCARE mission to perform on-board software maintenance operations depending on the size and the unit. The experienced gained from three in-flight software patches in both OBC and instrument have provided several lessons learned and a concrete workflow to decide a command distribution and command uplink strategy depending on specific constraints.

The trade-offs performed are applicable to any other missions under similar scenarios and could be used to optimise OBSM operations for LEO spacecrafts.

## Acknowledgements

ESA Spacecraft operations are most definitely a team effort requiring the support of multi-disciplinary expertise across ESA and industry and in the case of EC also JAXA. The authors gratefully acknowledge the excellent contributions of the wider EarthCARE teams in ESOC and ESRIN, the Project team at ESTEC, the Mission Manager and Airbus Defence & Space with all associated subcontractor experts to the EarthCARE mission and to the work behind this paper.

## References

- [1] Illingworth, A. J., H. W. Barker, A. Beljaars, M. Ceccaldi, H. Chepfer, N. Clerbaux, J. Cole, J. Delanoë, C. Domenech, D. P. Donovan, S. Fukuda, M. Hidakata, R. J. Hogan, A. Huenerbein, P. Kollias, T. Kubota, T. Nakajima, T. Y. Nakajima, T. Nishizawa, Y. Ohno, H. Okamoto, R. Oki, K. Sato, M. Satoh, M. W. Shephard, A. Velázquez-Blázquez, U. Wandinger, T. Wehr, and G.-J. van Zadelhoff. "The EarthCARE Satellite: The Next Step Forward in Global Measurements of Clouds, Aerosols, Precipitation, and Radiation", Bulletin of the American Meteorological Society 96, August 2015
- [2] Wehr, T., Kubota, T., Tzeremes, G., Wallace, K., Nakatsuka, H., Ohno, Y., Koopman, R., Rusli, S., Kikuchi, M., Eisinger, M., Tanaka, T., Taga, M., Deghaye, P., Tomita, E., and Bernaerts, D.: The EarthCARE mission – science and system overview, Atmos. Meas. Tech., 16, 3581–3608, 2023
- [3] European Cooperation for Space Standardization, Telemetry and telecommand packet utilization (ECSS-E-ST-70-41C), 2016.
- [4] Morales, J., "Rosetta On-Board Software: Design and Maintenance Concept for an Interplanetary Mission", Deutscher LuftundRaumfahrtkongress, Hamburg, 2002.

[5] Morales, J., C. Steiger, E. Montagnon, P. Ferri, “Rosetta On-Board Software Maintenance at the Control Centre: a Way to Reduce Operations Cost”, 5th Symposium on Reducing Cost in Ground Segment and Operations, ESOC, Darmstadt, June 2005.

[6] C. Steiger, G. Lautenschläger, A. Weigl, R. Eilenberger, P. Väisänen, E. Maestroni, A. da Costa, P. P. Emanuelli, “On-board software maintenance for GOCE, ESA’s gravity mission”, August 2010

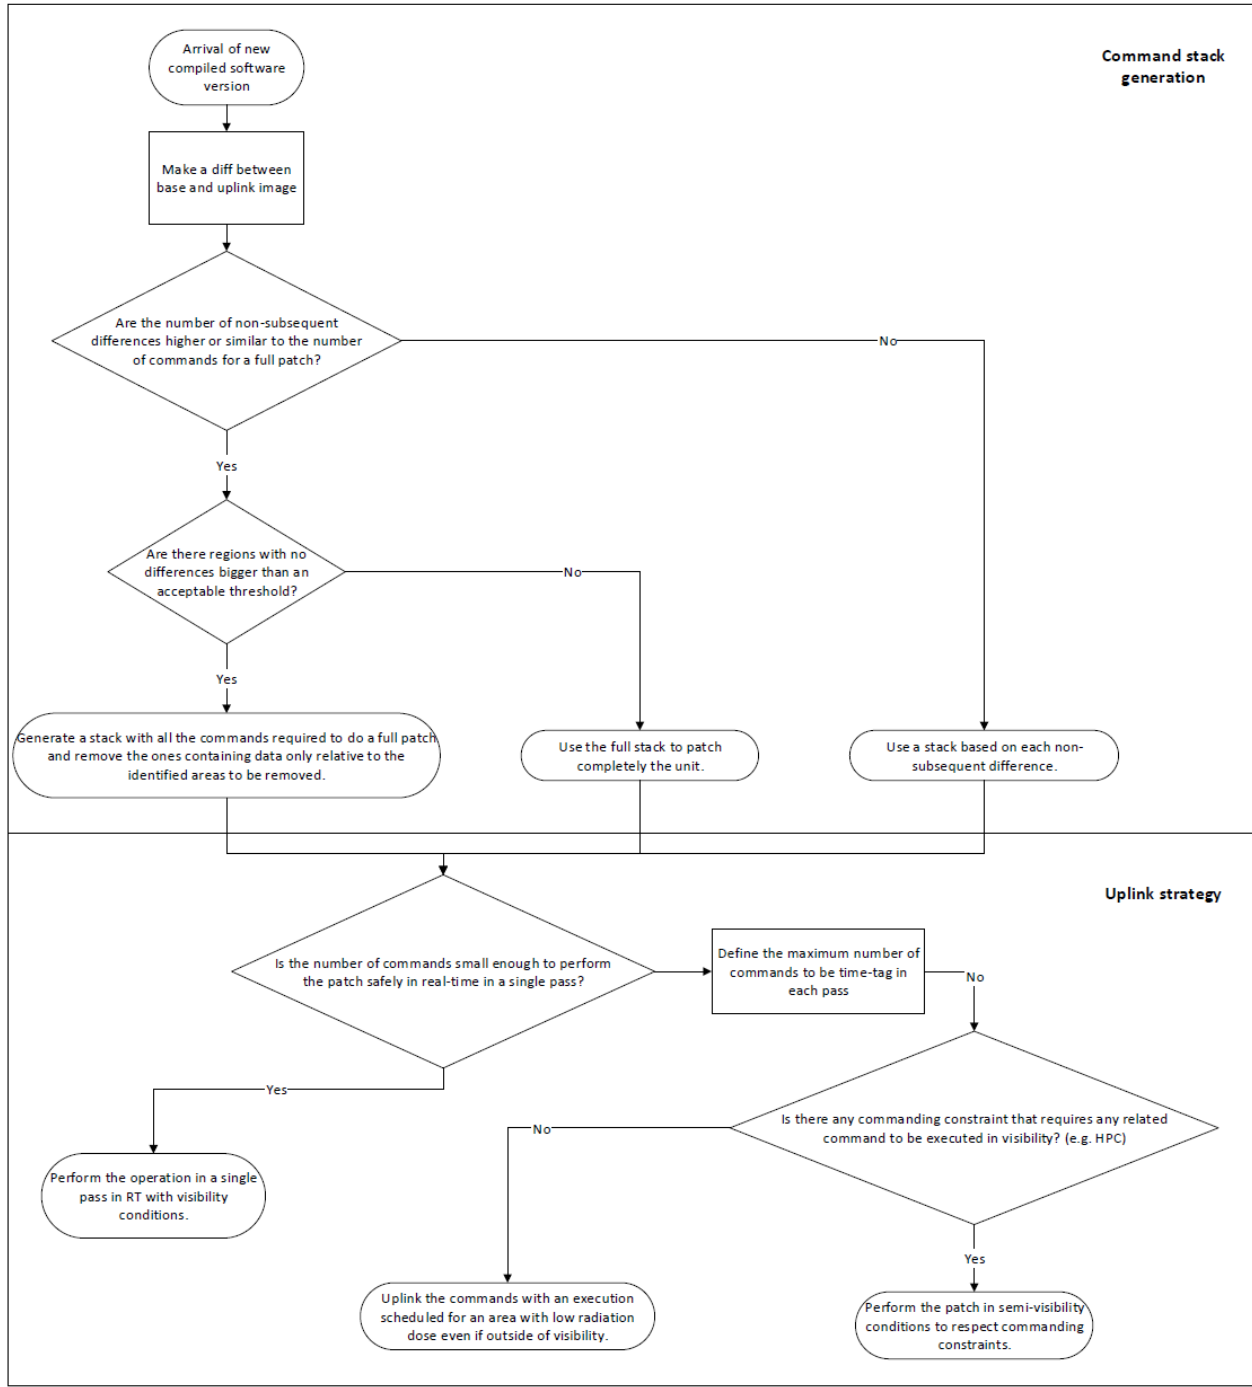


Figure 6 Workflow followed for the OBSM operations performed in EarthCARE during commissioning