Performance Evaluation of Arctic Weather Satellite

Ground Segment Development Plan

Authors

Adam Dybbroe, SMHI
Timo Ryyppö, FMI
Trygve Aspenes, MET
Panu Lathinen, FMI
Martin Raspaud, SMHI

Friday 17th February, 2023

Contents

Table 1: Document version control.

Version	Comment	Date
1.0	Initial version delivered for DDB review 7 Dec 2022	2022-11-30
1.1	Update solving the issues raised at the DDB review: • Added compliance statements concerning AWS acquisition for the four stations in section 2. • Add in section 3.4 specification of AWS test data required to be able to adequately test the AWS processors provided by ESA and the AWS industry. • Adding a brief discussion and assessment of the hierarchical schedule sharing approach and its vulnerability to contingency situations, in a new sub-section 2.3.3. • Update figure 7 and the caption text, reflecting that it is the raw data (not level-0) that are sent from Greenland to SMHI. Also made the share of responsibility between DMI (and the project) and EUMETSAT clear (EUMETSAT deliveres raw data to a DMI server and then on it is under responsibility of the project) in the text in section 2.	2023-02-17

1 Introduction and motivation

Under the Arctic Weather Satellite (AWS) industrial contract there are commitments only for one year of operation of the AWS prototype satellite, and the results of the AWS impact assessment (Task 5 of this project) using real data shall be presented at the Commissioning Results Review (CRR) just 9 months after launch. This means it is crucial to be able to provide an operational real-time flow of level-1 data as early as possible after launch. While it is expected that EUMETSAT will set up a regional EARS-like service in the future if the EPS-Sterna constellation materializes it is not expected that EUMETSAT will be ready to and choose to set up an EARS trial service from day-1 based on the AWS prototype. The AWS satellite launch, currently planned for Q1/Q2 2024 will be rather close in time with the Metop-SG-A1 launch (currently expected beginning of 2025) and EUMETSAT will naturally put highest priority to secure regional services on EPS-SG data.

It is for this purpose a Nordic AWS Direct Data Broadcast (DDB) ground segment is being setup as part of this ESA project. The data being acquired early after launch will be evaluated in an NWP setup using at least two of the domains stipulated in Figure 1, one Arctic (either the Arome-Arctic domain or one of the Greenland domains) and a Nordic domain (e.g. the MetCoOp domain).

This document describes the Nordic DDB ground segment in detail and how it is being setup to be ready to process and disseminate AWS satellite data from launch and onwards. Data will be made available for the project internally in real-time and to the wider public should anyone be interested getting AWS level-1 data over northern Europe, including Greenland and parts of the Arctic and the Nordic region.

Figure 1: Geographical outline of the Harmonie-Arome NWP model domains used now and in the near future within the Nordic Meteorological Institutes. MetCoOp and Arome-Arctic domains left, the Danish/Icelandic operational domains in the middle, and the planned UWC-west domains to the right.

2 Ground station network

The Nordic AWS DDB ground segment will use the three Direct Readout (DR) stations within the domain of the project partners, namely the station in Norrköping, Sweden (SMHI), Oslo, Norway (MET Norway) and Sodankylä, Finland (FMI).

In addition, data from the EUMETSAT EARS station at the DMI premises in Kangerlussuaq, Greenland, will be used. An agreement between this project and EUMETSAT has been reached such that EUMETSAT will provide all AWS data received in Kangerlussuaq in raw format (CCSDS packets) to the project in real-time. The raw data will be made available by EUMETSAT on a DMI server in Greenland and from there sent to DMI in Copenhagen. EUMETSAT has estimated that around 60 % of all AWS overpasses visible from the station will be received and sent on to the project.

An overview of the antenna coverage as seen from the Kangerlussuaq station and the three Nordic sites for a 600 km orbit compatible with AWS is shown in Figure 2

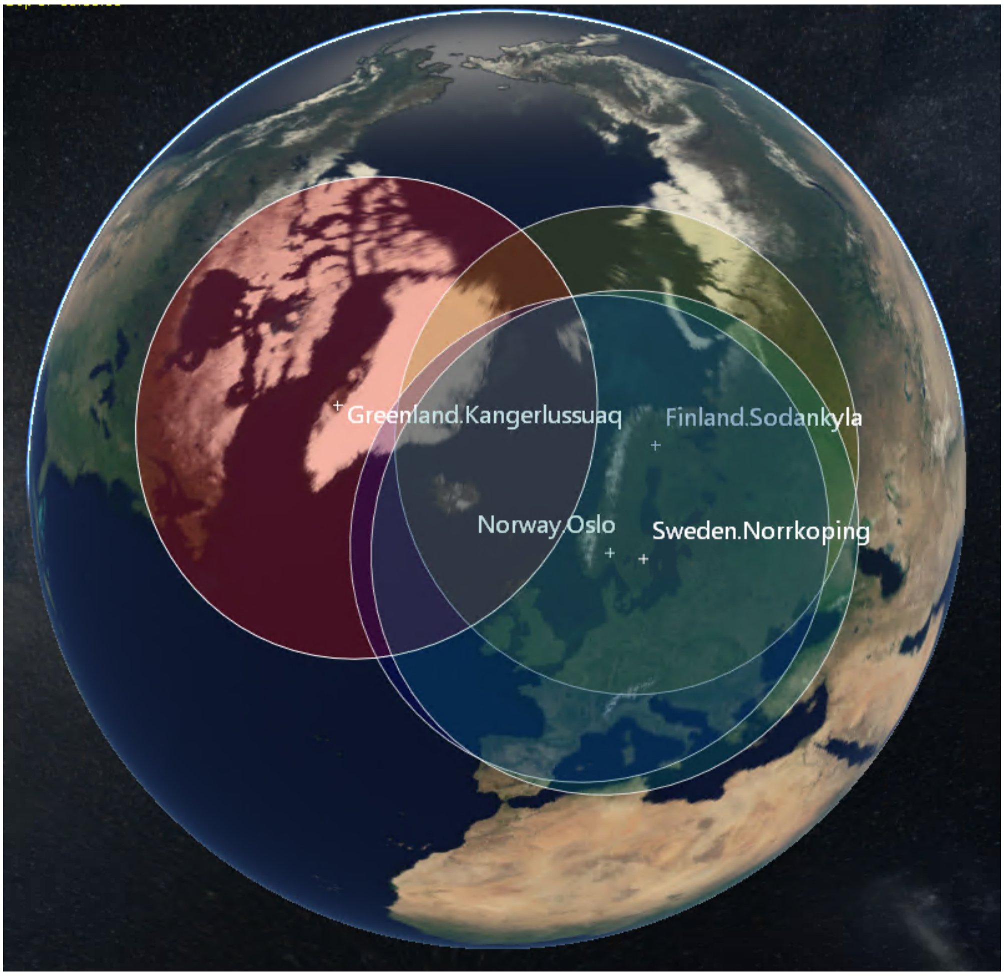

Figure 2: Antenna coverage for a 600 km polar orbit for a network of stations in Kangerlussuaq, Sodankylä, Oslo and Norrköping.

2.1 The stations and AWS acquisition readiness

All these sites currently have both X- and L-band reception capability, and they will be maintained and further developed at least until the end of this project.

All the four sites mentioned above today routinely acquire Metop and NOAA POES data in L-band.

Below a brief description of the four stations is provided, also in terms of their current AWS readiness and how full AWS readiness will be reached at launch.

2.1.1 Kangerlussuaq

The reception antennas (two Orbital X/L-band systems) and scheduling at the DMI premises in Kangerlussuaq, Greenland, are owned and operated solely by EUMETSAT and used for the various EARS services. This means that any scheduling priority (see Section 2.3 further down) for AWS at Kangerlussuaq and upgrades to the reception and processing has to be coordinated, agreed with, and performed by EUMETSAT.

Figure 3: The two X/L band antennas at Black Ridge, Kangerlussuaq, Greenland.

The reception system is Kongsberg's MEOS Polar $v5.3$. With the Space to Ground ICD ([1]) at hand and after a dialogue with EUMETSAT and the station provider, it has been concluded that no hardware upgrade of the Kangerlussuaq station is required. The compliance for acquisition of AWS data has been checked against the ICD ([1]) regarding the current G/T performances of the X/L-band antennas, and the station is today fully compliant. Only a software configuration change is needed to accommodate the reception and acquisition of AWS data.

2.1.2 Oslo

MET Norway is also in possession of an X/L-band Direct Readout (DR) station operated automatically 24/7. The DR station is a MEOS station from Kongsberg Spacetec. The station is recently upgraded to MEOS $v5.2.12$, which makes MET Norway ready to receive Direct Readout from EPS-SG-A/B. As the Kangerlussuaq station, the MET Norway station will also be able to acquire AWS/EPS-Sterna data after a small configuration update. The G/T at $1707 MHz$ is 12.8 dB/K, thus compliant according to the ICD (see table 6-9 in [1]).

Figure 4: The X/L band antenna in Oslo.

The station today receives and processes around 50 overpasses a day from various polar orbiting weather satellites. Including Metop-B/C, NOAA-18/19/20, Suomi NPP, AQUA, TERRA and Fengyun-3D. MET Norway has a long experience of real-time acquisition and processing of DR polar satellite data in an event-driven chain to ensure fast access to critical users. The level-1 to level-2 processors being operated currently include

CSPP (for VIIRS), SeaDAS (for MODIS), AAPP (for AVHRR, AMSU-A and MHS), FY3D (MERSI-2, MWHS), NWC/PPS (for all polar imagers), and NWC/GEO (for Meteosat SEVIRI).

2.1.3 Sodankylä

FMI operates the Arctic Space Centre (ARC) in Sodankylä which consists of four antennas, operations centre and processing and archiving infrastructure. FMI-ARC receives on a daily basis close to 70 satellite overpasses which data is used both in real-time operations (weather models and icebreaker assistance) and supporting FMI research activities (calibration and validation). FMI-ARC is Finland's Copernicus Collaborative Ground Station and a national data-hub to access Sentinel satellite data.

Figure 5: One of the two 3.7 m X/L band antennas in Sodankylä.

FMI has two 7.3 m antennas operating in X-band and two 3.7 m antennas operating in X/L-band. The antennas work in pairs that are connected with IF (Intermediate Frequency) matrix switches. This approach allows FMI to seamlessly change the acquisitions of different satellites between antennas and "overpopulate" the schedule with advanced conflict analysis process. The multi-antenna approach also provides redundancy, resilience and fault tolerance to operational ground station services. Like on Oslo the two L-band capable antennas are ready to acquire AWS data after a small configuration update. Recent measurements from 2021 of the G/T at the frequencies 1685.5 to 1705.5 (close to the planned AWS downlink frequency - see [1]) is between 11.3 and 11.5 dB/K, thus compliant according to the ICD (see table 6-9 in [1]).

At the moment FMI is actively receiving data from several, mostly weather related satellites. The list of satellites includes among others Suomi NPP, NOAA-20, EOS-Terra/Aqua satellites, Metop-B/C, Cosmo-SkyMed constellation and HaiYang-2B/C. Some of the received data, like VIIRS from Suomi-NPP and NOAA-20, are processed on-site to L1 and L2 products. In addition to processing the data for internal usage, FMI also offers ground station services. On daily basis FMI receives data from HaiYang satellites for EUMETSAT and distributes the raw data for processing and further dissemination through EUMETCast.

2.1.4 Norrköping

SMHI is currently operating an X/L-band Direct Readout (DR) station with a 2.8 m antenna dish. The system is installed in a Radome on top of one of the buildings of the SMHI headquarter in Norrköping and is provided and supported by CGI Germany (formerly SCISYS). The acquisition server is running `2met!` *Polar Acquisition* `V4.5`.

The station is currently routinely receiving and processing

Figure 6: The X/L band antenna in Norrköping captured just after installation and without the Radome, 2011

on average around 50 overpasses a day from various polar orbiting weather satellites. Currently data from the following missions are routinely acquired and processed to level-1 and level-2 products and imagery:

- JPSS satellites: NOAA-20 and Suomi-NPP
- EOS-Terra/Aqua
- NOAA-POES series (NOAA-18 and -19)
- Metop satellites (Metop-B and -C)

The 3rd party processors being operated include CSPP (for VIIRS), SeaDAS (for MODIS), AAPP (for AVHRR, AMSU-A and MHS), and NWC/PPS (for all polar imagers). In addition the Real-time Software Telemetry Processing System (RT-STPS) from NASA is run on the CGI 2met acquisition server in order to ingest raw telemetry data and producing products, including sorted Consultative Committee for Space Data Systems (CCSDS) packets and Virtual Channel Data Units (VCDUs).

The station was installed in November 2011. A procurement process has recently started in order to replace this station with a new one, capable of also receiving the data from the EPS-SG satellite series. The new station will replace the old one and be positioned within the existing Radome. The current time plan for this new station at SMHI is that it will become operational Q4/2023. The system will of course be AWS compliant against the ICD which has been provided to all bidders. The L-band G/T of the current system was measured to 9.25 dB/K at the On-Site Acceptance Testing during installation in 2011, and thus in that respect ready for AWS acquisition. Furthermore SMHI has been in dialogue with CGI about the AWS readiness of the existing system which will require a software upgrade only. Should the new system, in contrary to expectations not be ready by AWS launch, SMHI will together with CGI see to that the required software upgrade is in place prior to launch.

2.2 Inter-connectivity

The three Nordic stations are all connected via Internet, and each one are connected to their respective Nordic research and education networks under the NORDUnet umbrella (`https://nordu.net/`)[1].

2.2.1 The three Nordic nodes

For the AWS ground segment the raw AWS data from Greenland will be pulled from DMI to SMHI. The data transfer will use a simple FTP transfer. In order not to overload the connection by a constant polling very frequently, the polling frequency can be determined by the need of the NWP modelling. So, for instance frequent polling for new AWS data can start a few minutes before the assimilation cut-off time is reached. If the analysis time is 12

[1] SUNET in Sweden, FUNET in Finland and UNINETT in Norway.

UTC and the cut-off time is +1 h 15 min frequent polling could start at 13 UTC and end at 13:15. In between the analysis times, polling can still be regular but much less frequent, say every half an hour. At SMHI the Kangerlussuaq data will be put on disk for further processing.

There will be no direct flow of AWS data between the three Nordic nodes. The locally received AWS data, and the Greenland data processed at SMHI, will be put in an S3 bucket in a DMZ environment at each respective Institute. From there the data can be pulled by the wider public including the project, see Section 4 for details.

To have an optimized reception of AWS data over the ground segment there will be an exchange of reception AWS schedules between the individual stations. The schedules in `xml` format will be shared over internet on a regular basis. Only the planned AWS passes at each station will be shared, thus the reception schedules stripped from all other passes than the AWS ones. The exchange of schedules will use S3, as also used for the AWS data, see Section 3.3.

The schedule sharing aims to optimize the receptions to accommodate as many AWS passes as possible over the Nordic area while at the same time trying to keep the reception of other prioritised polar satellite passes as far as possible unaffected.

2.2.2 Connectivity to Greenland

Data connectivity between Kangerlussuaq, Greenland, and DMI in Copenhagen is constrained due to high pricing, resulting in a data line currently with a rather low capacity. The data connection has been used by DMI in the past to get satellite data and products over Greenland to the forecasting office in Copenhagen. However, this has now been terminated. Instead DMI gets the necessary data indirectly via the EARS service over EUMETCast. But, for the duration of this project DMI has kindly agreed to have the line in operation. It will be used to transfer the real-time raw (CCSDS packets) AWS data as well as weekly reception schedules in `xml` format. In agreement with EUMETSAT, the raw data are delivered from the antennas by EUMETSAT to a DMI server in Kangerlussuaq, and then further sent to Copenhagen and from there made available to SMHI on ftp. (See further in figure 7.)

2.3 Scheduling

The Arctic Weather Satellite (AWS) proto-flight model (PFM) is expected to fly in a low earth polar orbit at a height around 598 km above ground. The orbit is sun-synchronous and the spacecraft will circle the earth approximately 14 times a day. The goal of this project is naturally to acquire data from all passes visible from the ground stations (see Figure 2 and covering the area(s) of interest in this project as detailed in Figure 1.

However, there are some limitations and conflicts of interest that affects this ideal 100 % AWS PFM acquisition efficiency. First, as stated above the operation of the Kangerlussuaq station is outside the project, and exactly what AWS data being acquired there will be decided and scheduled by EUMETSAT. But, also as stated previously, EUMETSAT has agreed to provide around 60 % of all data visible from Kangerlussuaq to the project. Furthermore all three Nordic stations have many other commitments and priorities on what missions

Table 2: Current satellite acquisition priorities at Kangerlussuaq

Orbit	Satellite	Satellite Acquisition Priorities
Mid-morning	Metop-C	1
	Metop-B	4
Afternoon	NOAA-20 (JPSS-1)	5
	Suomi-NPP	2
	FY-3D	7
Early-morning	NOAA-18	6
	NOAA-19	3

to acquire, and due to their northerly positions many daily conflicts with more than one prioritised satellite within the local horizon at the same time. Both the stations in Oslo and Norrköping have only a single antenna to be considered for AWS acquisition, so there conflict resolutions are of particular importance. At Sodankylä there are four antennas of which two are L-band capable antennas and can be used for AWS. Regardless of FMI having two L-band antennas, the X-band missions can be acquired with all four antennas reducing the burden of L-band capable antennas.

2.3.1 Conflict resolution

When a Direct Readout station needs to choose which of two or more satellite overpasses within the local reception horizon to acquire at any given time there is a need for a clear and objective set of rules how to make this choice. At the EARS station in Kangerlussuaq, Greenland, the reception schedule is derived based on a simple satellite priority list, decided by EUMETSAT, see Table 2 depicting the current (as of spring 2022) relative priorities at the EARS stations. For example, according to this priority list, if a given Metop-C pass is conflicting with that of a Suomi-NPP pass, the Metop-C overpass will be scheduled for reception and the Suomi-NPP pass will then not be received.

In order to be able to decide a sustainable and reasonable scheduling strategy for the three Nordic stations that would at the same time maximise AWS reception and minimise the negative impact on reception of the other polar orbiters of critical interest, we have conducted a few schedule generation simulations. The simulations presented further below in Section 2.3.4 assume an AWS Local Time Descending Node (LTDN) of 10:30 UTC and is based on the satellites listed in Table 2, and using Two-Line Elements (TLEs) for these spacecrafts as of end of March 2022. The AWS TLEs were kindly provided by EUMETSAT.

2.3.2 Planned approach for optimized operational AWS scheduling

For the operational AWS ground segment and the four stations a simple hierarchical schedule sharing approach will be pursued. First, EUMETSAT schedules approximately 60 % of all AWS passes visible from Kangerlussuaq. AWS is here assigned the lowest priority. The Kangerlussuaq schedule is sent to DMI and shared to the three other stations. A Sodankylä schedule is generated using the knowledge of what AWS passes have already been scheduled for reception in Greenland. Then the Sodankylä schedule is shared with Norrköping.

From the list of AWS passes already scheduled at Kangerlussuaq and Sodankylä a schedule is generated for Norrköping, and finally Oslo will try take the possibly remaining AWS passes within the antenna horizon at Oslo, based on the prioritisation scheme applied at Oslo.

The AWS priority at Oslo will be higher than all other satellites except the VIIRS/ATMS satellites.

At Sodankylä the top priority is given to the JPSS missions (Suomi NPP, NOAA-20, NOAA-21) that is used in satellite based safety and security services. With the current launch schedule, AWS can be given high priority after the JPSS missions. Using the multi-antenna approach FMI is able to better secure the high priority of AWS and, in case of possible conflicts, modify the schedules of each antenna to better support AWS and other missions as well.

For Norrköping we plan for two options for AWS prioritisation:

1. AWS assigned highest priority among all polar orbiters
2. AWS assigned highest priority except all VIIRS/ATMS satellites (currently NOAA-20 and Suomi-NPP[2])

In the context of the *Nordsat roadmap towards operational data sharing* it is planned to have a data transfer of ATMS SDR data from Oslo to Norrköping in real time during 2023. Further it is also planned to have a real-time transfer of VIIRS SDR data from Oslo to Norrköping during the period when the SMHI station is being upgraded, see time plan in Section 5. VIIRS data are particular critical for SMHI, and so is the VIIRS timeliness. Among other things SMHI runs a real-time forest fire detection service based on local VIIRS Active Fires data for other authorities in Sweden. But if the VIIRS SDR data service from Oslo proves feasible (robust, reliable and with good timeliness) option 1 above can be a sustainable solution for both SMHI and the overall Nordic AWS ground segment. Therefore, we here consider option 1 above as the baseline and option 2 the fallback. The fallback option will be chosen in case the data sharing from Oslo should not prove feasible. Should option 1 be selected, option 2 can still be invoked temporarily should the transfer from Oslo become unreliable or during periods where local VIIRS data at SMHI is judged extra important.

2.3.3 Hierarchical schedule sharing and vulnerability to contingency situations

The hierarchical scheduling explained in 2.3.2 is somewhat vulnerable to situations where one station is not able to receive a committed AWS pass. This may happen if the station is suddenly down on short notice. In such a case the schedule will already have been shared with and considered by the other stations, and such an AWS pass may then be lost even if station specific conflict resolution would have allowed some of the other stations to receive it. However, there are reasons we do not expect this to be a major issue:

[2] As of today there are three JPSS satelites in orbit with the launch of JPSS-2 becoming NOAA-21. However, it is assumed that by the time of AWS there will only be two left to consider, namely NOAA-20 and NOAA-21.

- Sudden, unplanned outages of a receiving station is not something that happens very often. Mean Time Between Failure (MTBF) for each of the three Nordic stations are in the order of several months.

- Once it has been acknowledged that a station for some reason or another is down or incapable of receiving AWS, two things will be done: The station downstream in the hierarchy will be informed (allowing it to bypass the schedule sharing) and a new schedule will be calculated emptied for AWS passes. However, the time to react depends on the kind of service-level (*24/7* or *office hours only*) implemented for the operational reception may vary between stations. With the assumption of an *office hours only* type service-level the period where AWS passes might be lost due to the hierarchical approach will vary between half a day[3] to around 3 days (typical worst case scenario of incident happening on a Friday afternoon).

- The negative impact only happens when the committed but missed AWS pass in one station is conflicting with another prioritised pass at the next station in the chain.

The main objective of sharing the planned schedules among stations as described in 2.3.2 is to occasionally allow stations to plan scheduling the reception of passes also from other satellites (where they conflict with an AWS pass). This will have the effect of overall increasing the acquisition of other (non-AWS) passes at each station and reduce the redundancy of AWS reception. It is only in contingency situations as outlined above that this will be unfavourable from an AWS reception perspective.

Of course, if we experience repetitive unplanned incidents at one station not successfully receiving a committed AWS pass, that station will be decoupled from the schedule sharing.

2.3.4 Simulating the AWS acquisition efficiency from one possible LTDN

A rather high priority will be given to AWS at all stations except Kangerlussuaq as outlined above. Below we present the simulation results of the two options presented in Section 2.3.2 above, the *baseline* and the *fallback* approach. Here the simulations are using Pytroll-schedule which is being used in operation since several years in both Oslo and Norrköping. Pytroll-schedule allows for a more sophisticated conflict resolution between satellites compared to the simple prority-score based conflict resolution described in Section 2.3.1 and used at the EARS stations. Pytroll-schedule is using individual weights for each satellite (one for day and one for night) and tries to optimise the overall reception for a station in terms of coverage in time and space (using a pre-defined area of interest).

In Table 3 the summary results for AWS reception for all four stations are presented, simulating the reception of all satellites listed in Table 2 giving a very low weight to AWS in Greenland, and otherwise following the two scenarios introduced in Section 2.3.2. In Table 4 only results for the three Nordic stations are presented ignoring what might be acquired in Kangerlussuaq.

As can be seen the relative amount of data received in both scenarios are always smaller than the relative amount of passes received. This is because for a given AWS overpass only

[3] The tentative plan is to update the schedules twice a day

one station is receiving that particular pass. If AWS was the only satellite to acquire one would be able to get more data by stitching together the data from an overpass acquired at several stations. However, the majority of the data not acquired will be outside the area where the reception horizons overlap. And as seen from Figure 2 the reception horizons overlap over major parts of the NWP domains considered in this project (see Figure 1), except Greenland and the upper/northern third part of the Arome-Arctic area.

Table 3: Detailed simulation results of the two considered scenarios for AWS acquisition, the baseline and the fallback. In both the scenarios studied here a total of 142 passes were visible, with a total length of 125103 seconds (≈ 34.75 hours).

Scenario	Total number of passes received	Relative amount of passes received (%)	Time received (sec)	Relative amount of data received (%)
Baseline	118	83	72671	58
Fallback	101	71	61917	49.5

Table 4: Detailed simulation results of the two considered scenarios for AWS acquisition, the baseline and the fallback. Here, not considering Kangerlussuaq, only focusing on what the three Nordic stations may acquire. In both the scenarios studied here a total of 113 passes were visible, with a total length of 84127 seconds (≈ 23.37 hours).

Scenario	Total number of passes received	Relative amount of passes received (%)	Time received (sec)	Relative amount of data received (%)
Baseline	105	93	65130	77
Fallback	85	75	52487	62

As also LTDNs for three additional possible AWS orbital planes were provided by EUMETSAT for the same period we repeated the same simulations above for those as well. The LTDNs were 07:30, 09:30, and 13:30. All these three LTDNs resulted in almost the same statistics as presented in Tables 3 and 4, deviating only be a couple of percent.

3 Real-time acquisition, processing and dissemination

A high level illustration of the data acquisition, processing and dissemination is shown in Figure 7. The raw AWS data received in Kangerlussuaq are sent directly to DMI in Copenhagen without any further processing. The Kangerlussuaq data will be put on an ftp server at DMI, and data will be pulled from SMHI.

Each of the three Nordic nodes will do the acquisition and processing to level-1 of its own data. In addition the Greenland data will be processed to level-1 at SMHI. Data from each node will be uploaded to a node-specific S3 bucket. There will thus be three separate S3 buckets one for each of the three Nordic nodes, SMHI, FMI and MET. Each bucket will contain only the local data from that node/station, except the SMHI bucket which will also contain the data from Kangerlussuaq. Having three separate buckets with local data only instead of one bucket with all data from all stations (including Greenland) will provide some extra resilience for the user. If one bucket is inaccessible for some reason there will still be rlevant data available from the two other buckets, and only the data unique to the node where the respective bucket is down will be missed. If there should be overlap/redundancy in the data provided between stations, this resilience of several separate buckets over one central one becomes even more pronounced.

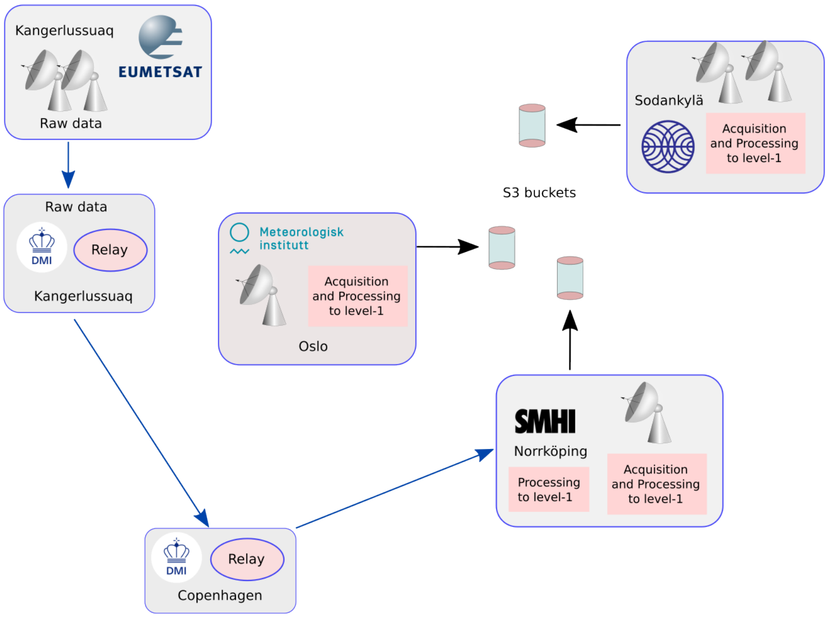

Figure 7: Sketch of the overall AWS data acquisition, processing and dissemination of the Nordic AWS ground segment, including the Greenland station. Raw AWS data received by the EUMETSAT antennas at Black Ridge outside Kangerlussuaq is sent to a DMI server in Kangerlussuaq and then relayed via DMI in Copenhagen to SMHI using ftp. The raw data are then processed at SMHI to level-1. Local data at the three Nordic stations are being processed to level-1. Level-1 data (including the Greenland data from SMHI) are finally uploaded to S3-buckets, one bucket for each of the three Nordic station.

3.1 Hardware

For both the off-line testing and the eventual real-time operation of the level-0 and level-1 processors delivered by the AWS Industrial consortium, it is sufficient to build upon the computational HW infrastructure already in place at the three Nordic sites and being used today for the processing of existing satellite missions.

This is under the assumption that the AWS level-0 to level-1c processing is not significantly more demanding in terms of disk space, CPU or RAM memory than processing e.g. ATMS or AMSU-A/MHS data. Therefore the ground segment development plan does not include any particular HW procurement in addition to what is already taken care of in the continuous procurement and upgrade plans at the individual reception locations.

3.2 Orchestration

The real-time operational setup will be leveraged from how the already existing level-1 processing (processing from level-0 to level-1) is done at the Nordic Institutes today; in general extensively built upon the Pytroll software tools (see [2]). An *AWS level-1 runner* will be developed similar to the corresponding *pytroll-runners* for the existing polar satellite data processing, e.g. MODIS, VIIRS, AVHRR and corresponding Passive Microwave (PMW) Sounder data.

The real-time level-1 processing chain will be developed to make deployment at each site as easy as possible. Containerized solutions (TBD) may be considered. The real-time level-1 processing chain will be deployed at each site before satellite launch. As the processing is very similar to the ATMS chain, see Annex 7, the concept is considered solid and well tested. However, it is the plan to have the AWS chain tested before launch with a pseudo live stream of level-0 test data simulating one satellite overpass. However, if such a live stream of level-0 test data is feasible with what will be delivered from ESA is TBD.

With the delivery of the final version of the level-1b and level-1c processors, these will be integrated and tested at one site (actual site TBD) in the production-ready level-1 processing setup, and an upgraded real-time level-1 processing system will then be delivered to the other sites for on-site testing prior to launch. The system will be delivered as a detailed set of installation instructions and/or as a ready-to-go container.

A detailed sketch of the general processing orchestration at one of the nodes is shown in Figure 8. It is here assumed that the level-0 processor will output data in granules of X minutes of observation data (actual value of X TBD). These segments/granules of level-0 data can then be gathered into a full swath before being ingested to the level-1b processing. It is yet TBD whether this gathering of granules will be invoked or not. Bypassing the granules gathering is easily accomplished via configuration. The message topics are configurable and determine which of the processors being triggered for what data. With shorter granules an overall lower data latency can be achieved. But it also requires the level-1b and level-1c software performs well on smaller chunks of data.

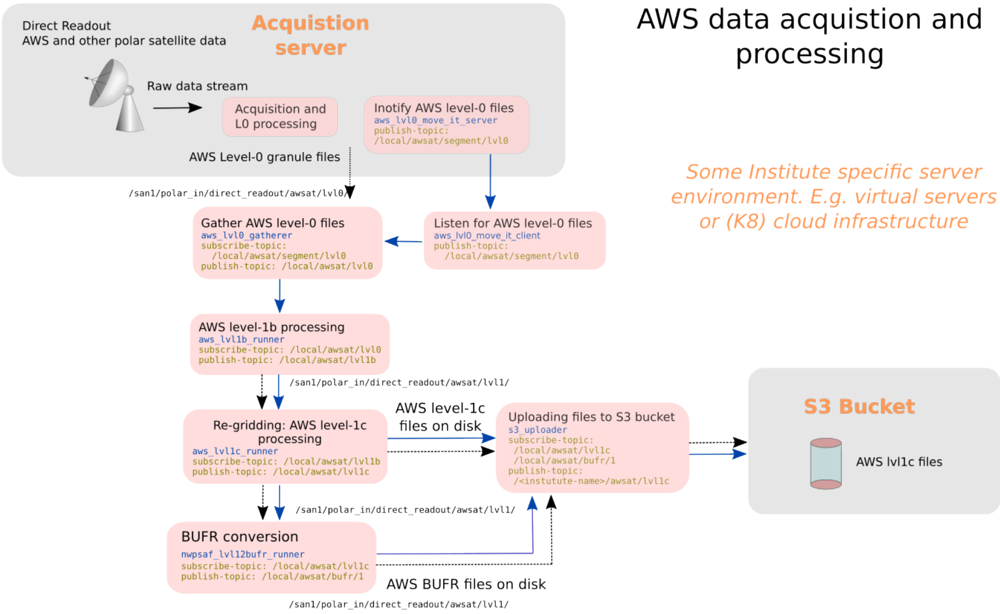

Figure 8: Sketch of the acquisition and processing of AWS data at one node (one of the three Nordic reception stations). Processing is done to level-1c, that is re-gridded AWS radiances and brightness temperarures in netCDF and BUFR format. Final level-1c data are being uploaded to an S3 bucket. The message topics used will likely vary between stations, here the illustration follows the one planned at SMHI.

The processing implementation is using various components from the Pytroll (see [2]) suite of open-source software packages. The concept is that of mini-services connected to each other via a subscriber/publisher interface. The most important Pytroll software packages being used are the following:

- Posttroll - for messaging
- Trollmoves - to move files from one destination to another
- Pytroll-collectors - e.g. to gather granules together
- Pyresample - library primarily for remapping geospatial remote sensing data, but have various support functions for spherical geometry calculations e.g. needed by Pytroll-collectors.
- Pytroll-schedule - for generating optimized local reception schedules
- Pyorbital - to make satellite orbital predictions based on TLE files and other astronomical calculations

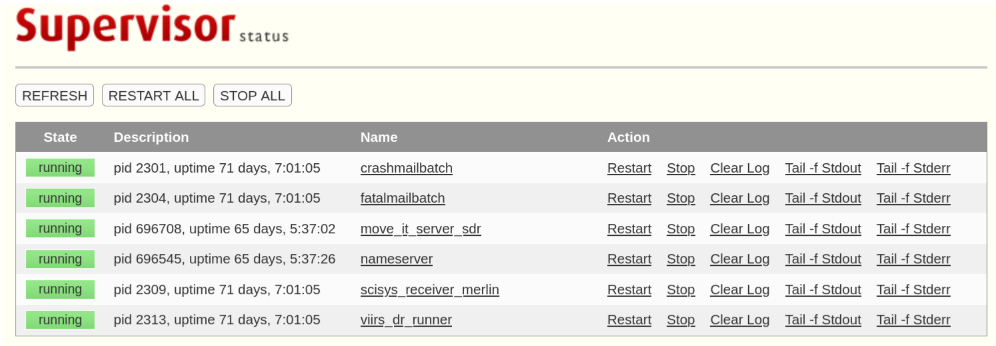

Figure 9: A screen capture of an example of a Supervisor GUI for an operational service in use at SMHI. This one is from the setup at SMHI for processing of locally received VIIRS RDR to SDR on one (VMWare) virtual server.

All processes are controlled and (re)started via Supervisor (http://supervisord.org/). An example processing chain using supervisor is shown in Figure 9. Supervisor takes care of keeping each mini-service alive. If a process crashes, for some reason, supervisor automatically restarts it again. If a restart fails after a number of automatic restarts (configurable) the process will turn red in the GUI, and an alert can be triggered in a Nagios® (`https://www.nagios.com/`) based monitoring service for instance.

The setup described above, using supervisor, e.g. in a virtual server environment as at SMHI, is how the operational satellite processing is done today at the three Nordic Institutes. However, there are work ongoing towards using cloud based processing. For instance at FMI work is progressing towards using OpenShift/Kubernetes for parts of the operational satellite processing. If such solutions are implemented for the AWS chains the use of Supervisor may become obsolete[4].

3.3 Dissemination

As outlined in Figure 8 the final level-1c files will be uploaded in real-time to an S3 bucket, one separate bucket for each of the three Nordic stations. The data files will contain either data from one local overpass of AWS data for the particular station or X minutes of data from one granule of a local overpass.

The data will be available for pulling from the S3 bucket for both the wider public and the project. There will be three distinct S3 buckets each holding the data from each of the three stations. In addition SMHI will have a bucket containing the data from the Kangerlussuaq station. To clearly identify where data were acquired a three letter identifier will be used consistent with how it is done for the EARS network today. The codes are outlined in Table 5.

[4] The (re)starting and control of the mini-services may then be handled by OpenShift/Kubernetes

Table 5: 3-letter station identifiers.

Station name	Station identifier
Kangerlussuaq	*kan*
Sodankylä	*sod*
Oslo	*osl*
Norrköping	*nor*

The WMO file naming convention will be followed as far as possible. Following the file naming of the EUMETSAT EARS-ATMS service this would be[5]:

```
W_XX-NORDSAT-<station>,<instrument>,DBNet+<satellite>+<station>_C_<CCCC-Code>_
<YYYYMMDDhhmmss>_<instrument>_<product-date-stamp>_<product-time-stamp>_<satellite>_
<orbit>_<service>_<station>_<product>_<extension1>.<extension2>
```

[5]The WMO CCCC code for the DB station in Sodankylä is EUMY, but the Oslo and Norrköping DB stations do not yet have any registered location indicators. Instead ESWI for SMHI, Norrköping, and ENMI for MET Norway, Oslo will be used here.

3.4 AWS processors testing and integration

ESA is responsible for the delivery of a set of AWS processors to process the raw downlinked DB data to level-1b, following the EPS-SG netCDF format. Two processor packages to fulfil this requirement will be delivered, a level-0 processor package and a level-1b processor package. In both cases the individual processor is considered to be a black box, producing an expected output on a specified input. Therefore, testing these entities require input test data and accompanying expected reference output data. These reference data will be used for verification against the locally (as a result of the on site testing) generated output.

Below follows a description of the test data required to adequately test and verify the local implementation of the ESA processors.

- Raw data
  - Raw data as required by the level-0 processor and in format following the space to ground ICD and as it would look like after downlink and reception.
  - The data should correspond to a duration of at least 3 minutes and preferably 10-12 minutes corresponding to a local overpass.
  - As a minimum the VCDU data should contain data in the Virtual Channel 0 (VC ID 0) (Real-time S/C HK telemetry data)
  - It would be beneficial to also have an additional test data set with data in the VC ID 0, 1 and 2, simulating the case where we will receive both direct real data and stored mission data
  - In order to achieve the required timeliness (6 minutes = goal; 15 minutes = breakthrough) it will be necessary to start process a local overpass while receiving. This means it will be required to be able to stream the data to a pre-processor (preferably part of the level-0 processor) that can chunk the data and provide smaller granules of data to the level-0 processor. Therefore it would be beneficial if ESA could provide the smallest feasible chunk of data for such a granule processing with the level-0 and -1 processors (it is assumed that there is a lower limit of how small a granule can be in order to achieve nominal data without requiring knowledge of more than one of the neighbouring granules each side of the granule being processed - for calibration purposes for instance).
- Level-0 data
  - Reference output data from the level-0 processor, corresponding to the raw test data provided in item 1 above. This will be used to verify the correct local installation and operation of the level-0 processor.
  - These level-0 data files will also be the ones that should be used to test the level-1b processor.
- Level-1b data

- Level-1b data files in netCDF corresponding to the level-0 files being input to the level-1b processor. These level-1b reference files will be used to compare width when installing and testing the level-1b processor locally. The data produced with the level-1 processor will be compared to the reference data, requiring no deviation in order to conclude a successful local level-1 processor testing.

- Level-1c data
  - The generation of level-1c data with the re-gridding level-1c processor is considered outside the scope of the AWS PFM mission. However, to test the generation of level-1c data locally it would be helpful to have a level-1c reference dataset corresponding to the test data above.

- TLE data
  - It will be beneficial to also have the (simulated) TLE data applicable to the data scenes provided (as outlined above in items 1-4).

Once the level-0 processor and accompanying test data are delivered the processor will be tested off-line at one or several of the sites in Sweden, Norway or Finland, TBD. Once successfully tested and verified at one site it will be delivered to all the three sites for local deployment to operations. Local deployment is planned just prior to the AWS launch. Included in the task to locally deploy the level-0 processor is also a mechanism to automatically dispatch the level-0 files (granules or full swath) from the acquisition server (usually a dedicated physical server provided by the station manufacturer/provider) to the level-1 processing environment.

When the Version-1 of the Level 1b and 1c processors and accompanying test data become available they will be tested off-line at one of the sites as outlined above for the level-0 processor. When successfully tested and verified off-line the development and implementation of a real-time operational setup can start.

3.5 Real-time process monitoring and up-time

The supervisor application is responsible for keeping the mini-services, outlined in Figure 8 up and running. If a process crashes it is automatically restarted. The number of restarts attempted can be configured. In addition all processes will be coupled to a central Nagios® type (or similar) monitoring service. At SMHI the op5 VMWare is used and at FMI `checkmk` (`https://checkmk.com/`) is used. At MET Norway Prometheus is used to collect and store metrics and alertmanager to check if metrics are within limit values. If not, an alert is triggered[6].

Such monitoring systems, as mentioned above, allows to either send a mail to the 3rd-line[7] support upon failure or couple the critical ones directly to the 24/7 operational moni-

[6] This alert can either be an email, webpages or SMS/automatic phone calls via pagerduty.

[7] At SMHI the 24/7 office is the 1st-line support, the 2nd line support is a back-officer from the IT department, on backup duty one week at a time typically. The 3rd-line support is the satellite processing experts responsible for the processing chain implementation.

toring desk at each of the three Nordic Institutes. The 24/7 operators may be able to act on the errors immediately or if not directly solvable defer to a 2nd-line support officer. It has to be noted while e.g. MET Norway does not have a 24/7 manned monitoring SMHI does, but still each system is classified whether a failure shall trigger an alert/alarm outside normal office hours, and very few systems have alerts that require action outside office hours also at SMHI. The AWS processing at SMHI is of a type which will normally not be classified as a system that requires action 24/7.

The response time from failure to a solution mitigation depends on the service level agreement decided for the actual service. For the AWS processing we propose a service level that secures immediate failure resolution only during office hours. Any failure resolution outside the normal office hours and which is beyond what the 1st line support can accommodate will be deferred to the next working day.

3.6 Archiving

All AWS data acquired during the first year in the Nordic AWS ground segment will be archived. Both data in level-0 and level-1 (level-1b or level1c, TBD) formats will be stored in the archive. Each station will archive its own data. The Kangerlussuaq data processed at SMHI will be archived at SMHI. There is currently no central archive planned holding all the data acquired in the Nordic AWS ground segment at one place. It is TBD whether all data (from all three nodes) will be gathered and de-duplicated at SMHI.

The purpose of the archive is to support scientific work within other WPs of this project mainly.

3.7 Points of Contact

In case a user discover problems accessing the data in one or several of the S3 buckets or experience other issues with the AWS data dissemination, it will be possible to reach out to a single point of contact via the following e-mail address aws-ground-segment@smhi.se. An e-mail sent to this address will then reach 3rd line support experts and relevant 1st line and 2nd line support destinations at the individual institutes. An issue raised via this channel will be taken care of within office hours only.

4 User side integration

A real-time user will be able to pull the most recent AWS level-1c data files directly from the three individual S3-buckets at MET, FMI and SMHI. Exactly how many hours/days of AWS data back in time will be available in the buckets is TBD. Functionality for doing this has recently been developed and added to the Trollmoves and Pytroll-collectors packages. Trollmoves has an s3downloader module that does the actual fetch and download of the files from S3. The Pytroll-collectors package provides a daemon process called s3stalker that with a certain frequency (configurable) go and fetch the url objects of the latest (configurable) files in the S3 bucket, and publish those urls for the s3downloader.

Thus the s3downloader listens to messages from the s3stalker, and download files to a configured destination. If the download is successful, i.e. the file exists on the local destination (e.g. a local disk drive) a posttroll message is published which may be used further downstream.

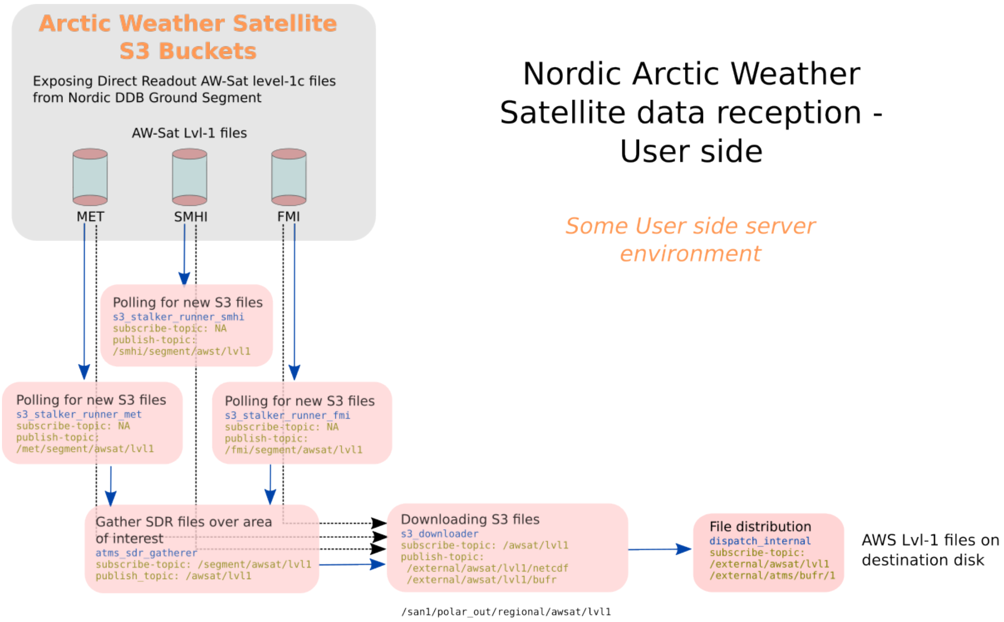

Figure 10: Example on a possible user side AWS data collection based on Pytroll. Blue arrow indicate communication flow, black dashed arrows indicate flow of data (files being moved/written on disk).

5 Ground segment setup - rough time plan

Below follow a tentative summary time plan setting up the ground segment prior to launch. Much of this work is carried out under the Project Work Package *1.2 Setting up Level-1b/c processing facilities and data distribution*, but there are also activities listed which are outside the scope of the project (like the ATMS data processing and sharing and the new station installation in Norrköping) but still necessary for the establishment of a successful AWS ground segment.

1. 2023/Q1: First tests and evaluation of the level-1b processing package
2. 2023/Q1: First setup of the local ATMS SDR processing and sharing between MET Norway and SMHI - Oslo data uploaded to S3 and SMHI local processing plus download of Oslo data.
3. 2023/Q2: Testing the level-1c processor with test data from the level-1b processor tests.
4. 2023/Q3: Final setup of the ATMS SDR data processing and sharing at all three nodes - no data from Kangerlussuaq.
5. 2023/Q3: First tests acquiring VIIRS SDR data in real-time from Oslo to Norrköping to be prepared for new antenna installation. This will also be a proof of concept for the routine backup of VIIRS data to SMHI during the one year AWS acquisition - following the baseline scenario outlined in Section 2.3.2.
6. 2023/Q4: New X/L-band antenna installation at SMHI, Norrköping.
7. 2023/Q4: First version of the AWS processing chain set up
8. 2024/Q1: Setting up the schedule sharing from Kangerlussuaq (here without AWS data of course).
9. 2024/Q1-Q2 Final setup of the AWS processing chain and dissemination.
10. Launch-2Weeks: All stations upgraded and prepared for AWS acquisition.
11. Launch
12. Launch+2Weeks: Final tuning and setup of the AWS processing chain and dissemination.
13. Launch+2Weeks: Hierarchical schedule sharing (only sharing the AWS passes to be acquired at each station) in place

6 Risks

A procurement process for a new direct readout station at SMHI, Norrkoping, has been started January 2022, and the current schedule aims at an installation during 2023/Q4, and thus approximately 6 months before the planned launch of the AWS PFM. The current station is compatible in terms of the L-band receiver, but a software upgrade is needed by the station provider CGI. Thus, there are two risks identified here, one is that the new station upgrade is delayed into the launch date of AWS, which could mean the station is not able to receive any polar passes for a period overlapping with the period when AWS is operational. The other risk is that the procurement process is halted for one reason or another and that the current station will have to operate the reception of AWS during the first year.

1. Minor: New SMHI station is not in place by launch
2. Minor: Current station is still the operational one at SMHI and configuration and software upgrade is necessary from the station provider - a special arrangement is needed for this as it is not currently within the contract.
3. The AWS ground segment is of course very dependent on the AWS processors delivered by the AWS Industrial Consortium via ESA and how well they integrate in the local infrastructure at hand for the three Nordic nodes. This is yet unknown, but as the delivery of the first version is available more than one year before launch we believe there should be time to correct any obstacles found during early testing.

List of Figures

List of Tables

Nomenclature

AWS Amazon Web Services

AWS Arctic Weather Satellite

CADU Channel Access Data Unit

CCSDS Consultative Committee for Space Data Systems

CRR Comissioning Results Review - for the overall AWS industrial project

CSPP Community Satellite Processing Package - provided my SSEC, Madison

DB Direct Broadcast

DDB Direct Data Broadcast

DMI Danish Meteorological Institute

DR Direct Readout

EARS EUMETSAT Advanced Re-transmission Service

EPS EUMETSAT Polar System

FMI Finnish Meteorological Institute

MET Norwegian Meteorological Institute

PMW Passive Microwave

RDR Raw Data Record - the term used for the JPSS missions, corresponding to level-0

RT-STPS Real-time Software Telemetry Processing System - provided by NASA

S3 Simple Storage Service - Amazon S3 is a service offered by Amazon Web Services (AWS) that provides object storage through a web service interface. But other Amazon S3 compatible object storage servers exist, e.g Minio

SDR Scientific Data Record - the term used for the JPSS missions, corresponding to level-1

SMHI Swedish Meteorological and Hydrological Institute

SSEC Space Science and Engineering Center - University of Wisconsin-Madison

VCDU Virtual Channel Data Units

References

[1] V. Garcia, "Arctic weather satelite direct data broadcast icd," 3 2022.

[2] M. Raspaud, D. Hoese, A. Dybbroe, P. Lahtinen, A. Devasthale, M. Itkin, U. Hamann, L. Ø. Rasmussen, E. S. Nielsen, T. Leppelt, A. Maul, C. Kliche, and H. Thorsteinsson, "PyTroll: An Open-Source, Community-Driven Python Framework to Process Earth Observation Satellite Data," *Bulletin of the American Meteorological Society*, vol. 99, pp. 1329–1336, July 2018.

7 Annex A - ATMS as a use case

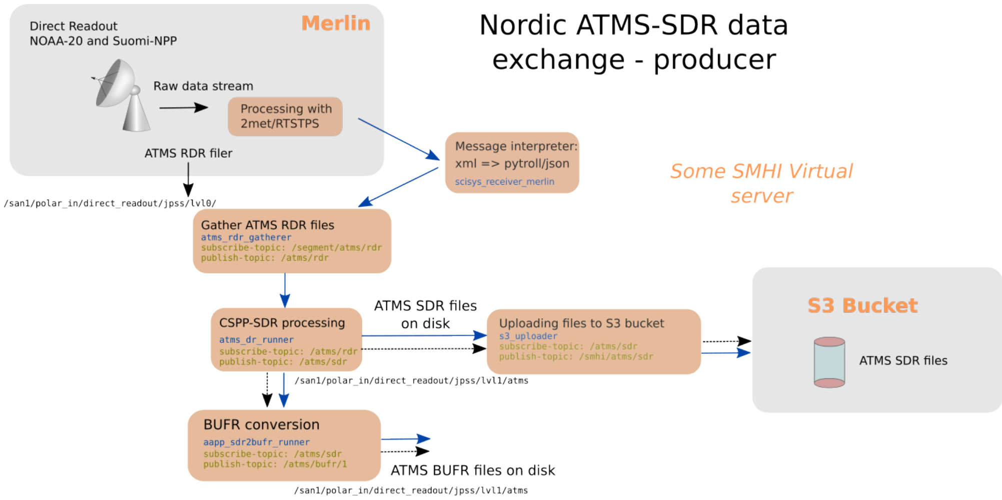

Figure 11: Sketch of of the planned data flow of the ATMS acquisition, processing and dissemination (upload to an S3 bucket) at each of the Nordic nodes. Here with tentative publish/subscribe topics at SMHI.

The ATMS processing chain is very similar to how the AWS/EPS-Sterna processing chain will be. It is therefore used as an example. Data is received from the Direct Broadcast from the satellite and raw data, i.e. CADU or other compatible format, is stored. For ATMS this is bundled with the other data from the JPSS platform so the RT-STPS tool is used to create ATMS Level-0 or RDR. RDR can either be one big file, or split into smaller segments to speed up processing. In the latter case you need to gather the RDR files before further processing. Seconds after the pass is complete the ATMS RDR file(s) are ready and moved to the processing server. Here CSPP is used to generate ATMS SDR (Level-1B) from RDR. These ATMS SDR files are then pushed to an S3 service where users can download the data. Also these ATMS SDRs are converted to BUFR using software from AAPP (provided by the NWPSAF), but there are also similar software in the CSPP package that can be used. These BUFR files are also uploaded to an S3 service and are suitable for data assimilation in NWP systems.

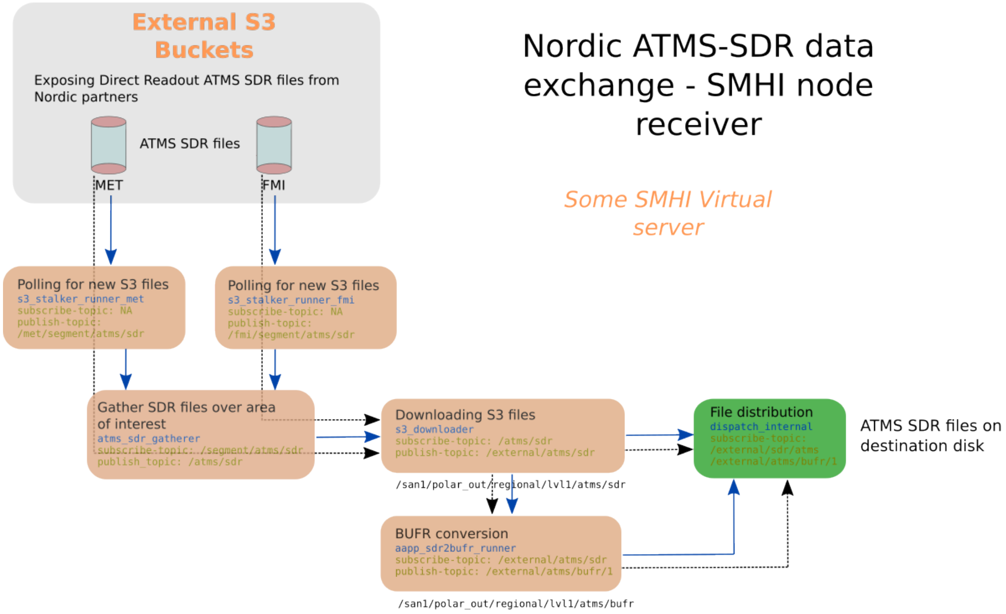

Figure 12: Sketch of the planned data flow at the SMHI node when fetching ATMS SDR data from the two other Nordic nodes and gathering only the relevant granules covering the area of interest (e.g. MetCoOp domain), BUFR conversion and internal dissemination. The setup will be similar at the other Nordic nodes.

8 Annex B - List of TBDs

1. The level-1b and level-1c processors will be integrated and tested at one site - actual site TBD.
2. The actual granule length in minutes or seconds is TBD. It will to some degree depend on the capability of the processors delivered by ESA.
3. Whether gathering of granules will be invoked or not before ingestion to the level-1 processing is TBD.
4. It is TBD whether all data (from all three nodes) will be gathered, de-duplicated and archived at SMHI.
5. Both level-0 and level-1 data acquired will be archived. But it is TBD which level-1 format (b or c) will be archived, probably level-1c though.
6. For packaging and sharing the full processing orchestration and AWS processor encapsulation among the nodes it is yet TBD whether containers will be used or a detailed set of installation instructions.
7. The naming of the level-1 files disseminated will follow the WMO file naming convention as far as possible. Precise file naming TBD.
8. Exactly how many hours/days of AWS data back in time will be available in the buckets is TBD.