

ASTER

Advanced Spaceborne Thermal Emission and Reflection Radiometer

Higher-Level Data Product Quality Assessment Plan

Version 2.0
October 18, 1999



Jet Propulsion Laboratory
California Institute of Technology
Pasadena, California

JPL D-13841

ASTER

Advanced Spaceborne Thermal Emission and Reflection Radiometer

Higher-Level Data Product Quality Assessment Plan

Prepared by: Craig Leff

Change Log

Date	Version Number	Reason for Change
February 5, 1998	1.0	Original document
October 18, 1999	2.0	<ul style="list-style-type: none">• Numerous changes in Appendix C, Tables 1-8, to correctly reflect delivered software.• Changes in Sections 2 and 5 to correctly reflect EDC DAAC roles and responsibilities.

TABLE OF CONTENTS

Change Log	i
1.0 Introduction	1
1.1 Document Overview	1
1.2 Overview of Basic QA Concepts	2
1.2.1 Bad and Suspect Data	2
1.2.2 QA Data Plane Concept	2
2.0 Roles and Responsibilities	4
2.1 DAAC	4
2.2 TLSCF	5
3.0 Product-Specific Topics	6
3.1 Product-Specific Attributes	6
3.2 QA Data Planes	6
3.2.1 First QA Data Plane	7
3.2.2 Second QA Data Plane	7
3.3 Replacement Values	8
3.3.1 Replacement Values for Level 1 Data	8
3.3.2 Replacement Values for Level 2 Data	8
4.0 Alerts	9
4.1 Alert Handling Scenario	9
5.0 Overall Flow	11
5.1 Automatic QA	11
5.2 Manual QA at the TLSCF - Investigative	11
5.3 Manual QA at the TLSCF - Routine	13
6.0 Trend Analysis	15
7.0 Data Volume and QA Workforce	16
7.1 Data Stream Sampling	16
7.2 Handling Bad Granules	17
APPENDIX A: Product-by-Product Alert Information	19
APPENDIX B: Product-by-Product Manual QA	27
APPENDIX C: ASTER QA Data Planes	30

1.0 Introduction

This document describes the quality assessment (QA) processes and components for the Advanced Spaceborne Thermal Emission and Reflection Radiometer (ASTER) Product Generation System (PGS) Version 2 and the EOSDIS Core System (ECS) Version 2. This document may be revised in the future based on the evolution of the PGS and ECS. It is intended for use by ASTER, DAAC and ECS QA personnel, by ASTER operations planning personnel, and by potential users of ASTER data products, including the ASTER Science Team (AST), Interdisciplinary Science (IDS) team members, and the science community at large.

1.1 Document Overview

Section 1.2 provides an overview of ASTER QA processing. Basic ASTER QA concepts are introduced there so that the casual reader gains a passing familiarity with the processes discussed in greater detail in the later sections of this document.

The subsequent sections of this document will discuss the details of the implementation of the QA concepts.

Sections 2 through 4 provide background information about the QA roles of the Land Processes Distributed Active Archive Center (hereafter referred to as the DAAC) and the ASTER Team Leader Science Computing Facility (TLSCF) (Section 2), the structure of the QA-related portions of ASTER data products (Section 3), and the handling of QA alerts, which trigger ASTER QA operations (Section 4). Supplemental details for all three sections are provided in the Appendices, which are described below.

Section 5 describes the end-to-end flow of ASTER QA operations, from the automatic checking of data products by production software using science-prescribed range limits to the visual checking of data products by members of the QA team to problem correction and documentation.

Section 6 describes trend analysis based on QA statistics. QA trend analysis provides insight into software and instrument performance and has implications for the scientific research of the ASTER Science Team (AST).

Section 7 provides estimates of both QA data volume and the TLSCF workforce resources needed to handle that volume.

Appendices A, B, and C provide additional details of the product-specific QA information: alerts, manual QA, and QA data planes.

1.2 Overview of Basic QA Concepts

ASTER QA processing can be divided into two types. The first type of QA is performed within the product generation software. In the course of data product generation, QA parameters are calculated and QA results are reported. Because this QA information is calculated and collected without user intervention, it is referred to as automatic QA.

During automatic QA, a variety of summary statistics are calculated that provide insight into product quality. If these summary statistics indicate a possible quality problem, an alert is raised, the alert is saved for future reference, the product is flagged as bad, and QA personnel are notified that manual QA is required for that product.

Manual QA is the second type of ASTER QA and is performed by a human operator. If automatic QA indicates that a product is bad or suspect, the DAAC first checks to see if there was some type of DAAC operational problem. If so, the problem is corrected and the product is re-generated. If no operational problem is found, manual QA is performed at the ASTER Team Leader Science Computing Facility (TLSCF) to determine the problem and whether it can be corrected. During manual QA the archived alerts may be used as ancillary information, along with other metadata. Identification and correction activities conducted within investigative QA are also recorded for future reference. Once these activities have been completed, the data product in question is marked as either good or bad.

There is also a second form of manual QA at the TLSCF, known as "routine QA." The normal data production stream will be sampled on a daily basis to provide an additional check on data product quality. Up to 10% of the daily granule production will be examined, with sampling criteria supplied by the ASTER Science Team.

Granules are formally defined as the smallest aggregation of data that is independently managed (i.e., described, inventoried, retrievable.) In practical terms for ASTER users, a granule is one scene of one data product (e.g., one particular example of the Temperature Emissivity Separation data product.)

1.2.1 Bad and Suspect Data

Bad and suspect data within a data product will not be replaced with cosmetic or marker values. The AST decided not to replace bad values because bad data may not be bad for all users or may still contain some scientifically useful information. In addition, bad data will be needed to understand why the pixels are bad, possibly leading to algorithm improvements.

1.2.2 QA Data Plane Concept

Each ASTER data product will have a minimum of one and a maximum of two data planes containing pixel-level QA information. Data planes are the logical representation of the data for each ASTER band. The QA data planes are created in addition to the science data planes and map to the science data on a pixel-by-pixel basis.

The first QA plane has a common structure for all data products (except DEM, which was developed by an independent contractor.) Each pixel will be marked as Good, Bad or Suspect, and its cloudiness noted. The second QA data plane is optional and its structure varies according to data product.

A data plane approach was chosen because it allows for easier graphical display of bad data (as well as other QA information), making it easier for a user to see which parts of the scene were most affected. It is also part of, and consistent with, the more general approach of storing QA information in a standard way for all ASTER products.

2.0 Roles and Responsibilities

2.1 DAAC

DAAC responsibilities for ASTER QA include:

- Initial evaluation of PGE failures to evaluate possible causes.

A PGE is a Product Generation Executable, a computer process used to generate ASTER standard products. The DAAC Science Support Group (science data specialists and Science Software Integration and Test engineers) will investigate problems with the PGEs, both performance problems and failures. This investigation will include the ECS environment in which the PGE is running operationally and the input files being used for operational processing. Problems include system failures and problems undetected in integration testing. These problems might include problems with the PGEs themselves or the ECS code in which the PGE is running such as the Toolkit or the PDPS subsystem.

- For problems deemed not caused by operational aspects of processing, collecting or identifying log files and granules necessary for TLSCF personnel to continue troubleshooting
- Supporting ASTER quality problems reported by users via the DAAC's troubleticketing system and User Services and coordinating these problems with the TLSCF.
- Ensuring the integrity of the data products and metadata, i.e., that data are not corrupted in the transfer, retrieval, or archival processes.

This shall be performed by the SSI&T process and system testing. Active checks on all transfers, retrievals or archivals will not be performed as a part of standard operations. The integrity of data products will be ensured through regular operation and monitoring of the ECS system and through troubleshooting of user reported distribution and data problems. Troubleshooting of data problems reported by users shall be limited to those associated with system failures such as missing or corrupted data files, bad tapes, etc. Data quality problems will be reported to the TLSCF as stated above.

- Supporting ECS subscription services so that AST personnel may automatically be notified of ASTER product generation and readiness for Quality Assessment.

Subscriptions permit users to register their interest in changes to and events associated with data products. Current functionality requires that the DAAC submit subscriptions on behalf of the interested customer. Subscriptions can be submitted for data of a given datatype being inserted into the ECS archive. Subscriptions have the capability to trigger e-mail notification of the event and also delivery of the data via

either an ftp push or 8mm tape. This trigger will occur with each archive insert of the requested datatype unless the subscription is qualified. Qualifiers are currently limited to the date and time of data acquisition.

2.2 TLSCF

TLSCF responsibilities for ASTER QA include:

- Ensuring the scientific integrity of ASTER data products. While granules may arrive from the DAAC with no operational errors, they are evaluated for acceptance by the AST based on both qualitative and quantitative examinations.
- Investigating possible science-related problems. By determining the cause of the out-of-range QA statistics, TLSCF personnel may lay the groundwork for solving and correcting these problems.
- Supporting manual investigative QA by facilitating the resolution of problems found during automatic QA.
- Performing manual routine QA.
- Interacting with DAAC personnel. Evaluation of data products and investigation of problems may require setting subscriptions, retrieving granules from DAAC servers, and working with DAAC personnel to retrieve and understand DAAC processing logs.
- Interacting with algorithm developers from the AST. Evaluation of data products and investigation of problems may also require discussions with the algorithm developers to understand an algorithm's behavior or to help the developer formulate changes to the algorithm to prevent future occurrences of bad data.
- Interacting with the instrument team and the AST regarding instrument-related problems. ASTER QA may reveal problems related to instrument behavior. Any such problems discovered would require discussions with the instrument team and the AST to help characterize the nature of the problem and its potential effects on the operation of the instrument and on data processing.

3.0 Product-Specific Topics

This section describes the product components that serve as containers for ASTER QA information, as well as the contents of those containers. ASTER scene-level QA information appears in the data product metadata, and ASTER has chosen to use product-specific attributes (PSAs) as the containers for those metadata. ASTER pixel-level QA information is collected during data product generation, and that information is stored in QA data planes, which are linked to the science data planes for each product.

3.1 Product-Specific Attributes

Product-specific attributes (PSAs) provide science teams with the ability to define metadata and to assign a value or values to that metadata. For ASTER data products, PSAs are used for the storage of alert-related information. QA-relevant statistics are calculated at the granule level within the ASTER Level 2 PGEs, and each of these statistics is stored in a PSA. QA PSAs are then used in determining alerts. An alert occurs when any of these QA statistics is not within their pre-determined ranges. Alerts will be discussed in more detail in later sections of this document.

Tables 1 and 2 in Appendix A describe all the alert-relevant PSAs. Table 1 describes the summary statistics, which are used to trigger alerts, and Table 2 describes the alerts themselves. All the summary statistics begin with the string "QAStat" and all the actual alert statistics begin with the string "QAAlert." Both the summary statistics and the alerts are listed on a per-product basis, along with a description and whether the alert is critical or non-critical.

3.2 QA Data Planes

In addition to science data, each ASTER data product includes one or more additional data planes containing pixel-level QA information. All ASTER data products (except DEMs) have one QA data plane; many of the products have one additional QA data plane. A data plane approach was chosen because it allows for simple graphical display of bad data (as well as other QA information), making it easier for a user to see which parts of the scene were most affected. It is also part of, and consistent with, the more general approach of storing QA information in a standard way for all ASTER products. This standard approach will make it easier for users to access ASTER QA information, as well as providing an efficient mechanism for storing this information on a pixel-by-pixel basis.

Shared QA information (i.e., information common to all data products) is stored in the first data plane. Any remaining space is used for product-unique information. Details of the first QA data plane are provided in Section 3.2.1 and in Appendix C.

The second QA data plane, when used by a particular product, contains additional product-specific QA information. More information about the second QA data plane is provided in Section 3.2.2 and in Appendix C.

3.2.1 First QA Data Plane

The first QA data plane is present in all ASTER products. It is composed of three fields:

1) The good/bad/suspect pixel code field, is four bits in length and occupies the most significant bits, as shown in Figure 1. Bit patterns for these four bits are used to specify categories of Good, Bad and Suspect pixels, based on information from developers.




Figure 1: The structure of the first QA data plane

2) The cloud status field is two bits in length and occupies the next two bits. These bits specify whether a pixel is cloud-free, adjacent to a cloud, covered by a thick cloud, or covered by a thin cloud. All of the ASTER data products use the cloud status field except the Polar Surface and Cloud Classification data product, which produces its own cloud information and thus does not use these two bits. The Polar Surface and Cloud Classification data product is a post-launch product and is described in more detail in Appendix B.

3) The adjacency distance pixel code field uses the last two bits and will provide estimates of the distance from the center of the pixel to the influencing cloud(s).

The pixel codes for all three fields are defined in Appendix C.

3.2.2 Second QA Data Plane

The second QA data plane contains product-specific QA information on a pixel-by-pixel basis. The size of the second data plane may be 8, 16, or 24 bits, depending on the needs of the algorithm developer. The information in the second QA data plane includes such information as:

- Computational Uncertainty
- Computational Path (e.g., the number of iterations of the algorithm until convergence was achieved or until processing was completed.)

Appendix C provides the information contained in the second QA data plane for each ASTER data product which uses it.

3.3 Replacement Values

3.3.1 Replacement Values for Level 1 Data

The Japanese ground data system performs Level 1A and 1B processing and is responsible for Level 1 QA and the interpolation of replacement values. While this document addresses higher-level data products (Level 2 and greater), Level 1 replacement values are briefly discussed here because of their fundamental influence on higher-level QA. Additional details can be found in the Level-1 Data Product Specification Document.

Data that are bad or missing in the Level 1A image will be replaced with interpolated values. This will minimize the diffusion of the bad data into the surrounding pixels which occurs as neighboring pixels are mixed during the resampling process. No additional interpolation is done in Level 1B processing.

Location information for interpolated pixels is documented in the header of the Level 1A data products. This header information is "carried through" into Level 1B and Level 2 data product headers and thus remains available to users.

3.3.2 Replacement Values for Level 2 Data

For Level 2 data, bad or missing data will not be replaced with interpolated values, and will be stored unchanged. If no value is calculated (e.g., due to an algorithm failure), it is set to the minimum value that can be represented in that data type (e.g., -32768 for a two-byte integer.) Similarly, if the calculated value is too large in magnitude to represent in the data type available, it is set to the maximum value that can be represented in that data type (e.g., +32768 for a two-byte integer.)

This approach was chosen by the AST for two reasons. First, bad data may not actually be bad, or may still be scientifically useful and so should not be replaced. The bad data may also be the only ASTER data collected over a certain location. Secondly, the bad data will be needed "intact" to understand why it is bad, possibly leading to algorithm improvements.

4.0 Alerts

In the course of Level 2 data product generation, QA-relevant summary statistics are calculated within the Level 2 PGEs for each granule. These summary statistics include values such as the percent of bad pixels in a scene. When any of these summary statistics is outside their assigned ranges, an "alert" is raised, causing alert information to be sent to an Alert Log, which is permanently archived at the DAAC for future reference. There will be one Alert Log for each ASTER data product. Alert information is also written to the data product header for the end-user's reference.

This following section details how alerts are calculated, collected and archived for ASTER data products.

4.1 Alert Handling Scenario

The scenario for calculating, storing and evaluating summary statistics and collecting alerts in the Alert Log is presented in the following steps. It is important to note that the terms "permanent" and "temporary" are used generically to refer to a saved file and to an interim file, respectively, and do not imply the ECS-specific usages of these terms.

- 1) During product generation, the value for each summary statistic is calculated. Each summary statistic is a Product-Specific Attribute and is reported as granule-level metadata, regardless of its value.
- 2) Each summary statistic is compared to its valid range. If the statistic is outside the valid range then the associated alert is triggered. If a "critical alert" occurs, the granule being processed fails automatic QA, the granule is designated for manual QA, and the alert is archived. "Non-critical alerts" are archived, but no specific QA action is prescribed. Valid ranges for summary statistics and the critical and non-critical alerts for each product are defined by the algorithm developers and stored in a look-up table, where they are adjustable, as necessary.
- 3) Each product will have a PSA called QASummaryofAlerts. This attribute consists of a text field containing a table summarizing all the alerts for the granule. Each time an alert is triggered another entry is made in the table. An example of this table is shown in detail in Appendix A.
- 4) The table stored in QASummaryofAlerts is written to two places:
 - a) The product header
 - b) In a temporary file associated with the current instance of the PGE
- 5) The number of critical alerts and the number of non-critical alerts are stored as metadata in the Product-Specific Attributes called QACritAlertsCnt and QANonCritAlertsCnt.

6) A temporary alert file is created by the PGE whenever one or more alerts are triggered during generation of a granule. Later, these temporary files are turned into permanent files which are archived. When this file is opened a standard set of information is first written to it to identify the source of the alerts. This standard set is TBD but will contain items such as:

- Timestamp
- Product name
- PGE Name
- Algorithm version
- Software version
- Granule identifier

After this header section, the table stored in QASummaryofAlerts appears. If no alerts are generated, this file is not created.

7) Periodically (TBR, but approximately weekly) a process is automatically started which concatenates all the temporary alert files into a single file (an ASTER alert granule called the Alert Log) and then archives that file. The temporary files are either deleted or set to expire after an appropriate period.

5.0 Overall Flow

Figure 1 provides a graphical overview of the ASTER QA scenario. Each step in the scenario has been numbered, and these numbers will be referred to as a guide throughout this section. Beginning at the upper left with Step 1, data product generation is requested via a Data Processing Request (DPR). Depending on the data product requested, the appropriate Product Generation Executable (PGE) is initiated to produce the product. Level 1B data are utilized by the PGEs to create the higher-level data products. However, Level 1B processing is not addressed in this document and occurs prior to Step 1 in Figure 1.

5.1 Automatic QA

In the course of data production, QA-relevant statistics for the granule are calculated within the PGEs (Step 2). Alerts encountered at this step are stored in a product-specific attribute (PSA), a metadata type that provides the means to define and store instrument-team-defined metadata information.

As discussed in Section 4, an alert is raised when any of the summary statistics is not within their AST-determined ranges. Alert information is sent to a temporary file via a PGE call (Step 3). Once per week, these temporary files are concatenated to create an alert granule known as the Alert Log, which is permanently archived for future reference by DAAC, TLSCF, and AST QA personnel. Alert information is also written to the data product header.

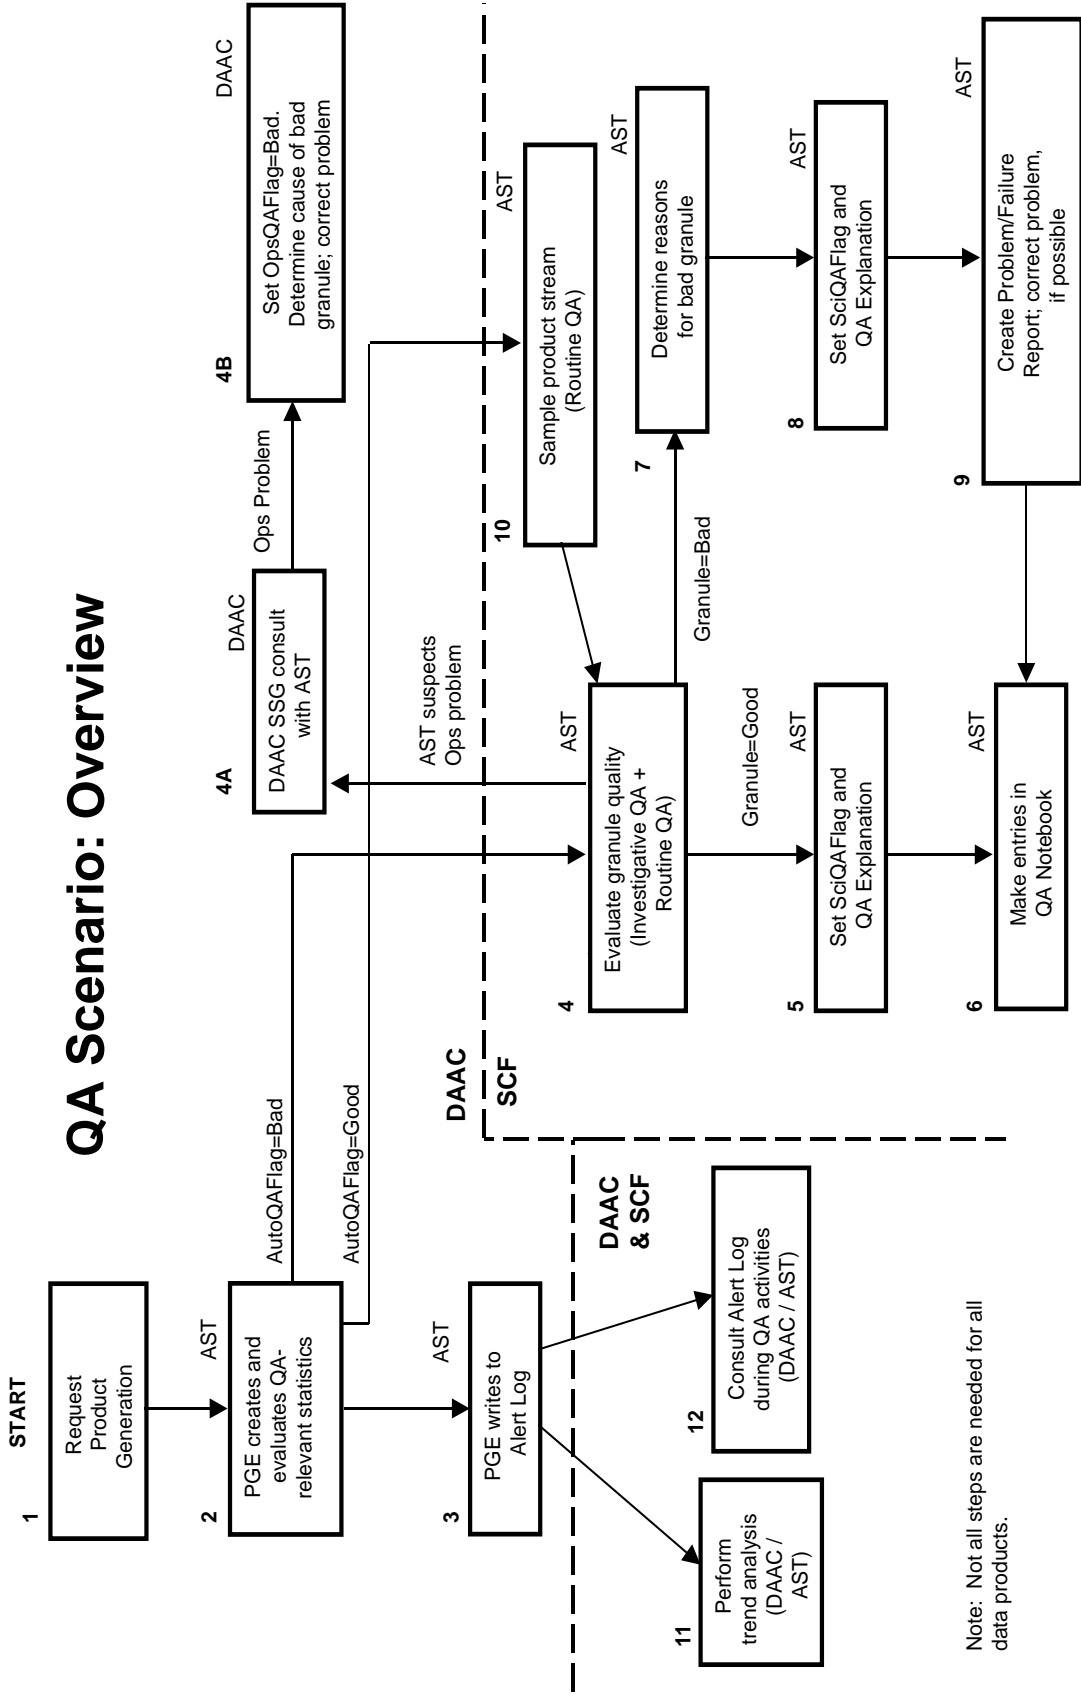
If the product passes automatic QA, the PGE sets the AutoQAFlag to Good, updates the header, and proceeds to the next data product in the queue.

5.2 Manual QA at the TLSCF - Investigative

Investigative QA refers to QA that is performed manually on products that have failed AutoQA (i.e., AutoQAFlag=Bad.) TLSCF personnel retrieve these granules from the DAAC and evaluate data quality as part of TLSCF investigative QA. During investigative QA, TLSCF personnel review critical alerts to see why the granule failed automatic QA and review the granule visually (Step 4). There is no set formula for determining the causes of granule failure, and problem solving is handled on a granule-by-granule basis. In addition to the experience of the TLSCF personnel, interaction with AST members, instrument team members, algorithm developers, and PGS software developers may be necessary to complete TLSCF investigative QA. During this process, DAAC personnel (the DAAC Science Support Group) are available for consultation if it is thought that an operational problem may be involved (Step 4A). If it is determined that it was an operational problem, the DAAC will set the OpsQAFlag to Bad and the product will be regenerated (Step 4B).

Once the causes have been identified, the Science QA Explanation field is updated using the ECS-provided Metadata Update Tool to explain why particular values were set for

QA Scenario: Overview



Note: Not all steps are needed for all data products.

Figure 1

the SciQAFlag metadata (Step 8). It is important to note that regardless of whether the Science QA Flag (SciQAFlag) is eventually set to Good or Bad, the AutoQAFlag will remain set to Bad for those granules where TLSCF investigative QA is necessary.

A description of each step performed during investigative QA will be captured in an entry to the QA Notebook (QANB), a file which serves to collect specific QA activities and comments in much the same way an experimental scientist would collect information in a laboratory notebook in their work area (Step 6). For example, if a batch of products was created using the wrong calibration files, and reprocessing was requested to correct the problem, QA personnel would compile this information to maintain a record of exactly what they did.

The purpose of the QANB is to capture exactly what was done during investigative QA so that it can be re-visited if necessary, and to provide a sequential summary of QA problems that may be useful for trend analysis (Step 11). The QANB is an ASCII file produced by TLSCF QA personnel and managed and maintained at the TLSCF. However, the QA Notebook will be available to all QA and AST personnel. Note that the QANB is only one component of QA trend analysis, which is discussed in more detail in Section 6.0.

When needed for TLSCF investigative QA, the Alert Log is retrieved from a designated directory on a designated machine at the DAAC (Step 12).

Problems discovered during TLSCF investigative QA will be captured in a Problem/Failure Report (PFR), which allows TLSCF QA personnel to formally document those problems (Step 9). The PFR may lead directly to a DPR to create a corrected granule or to the implementation of operational changes to prevent future occurrences of the problem.

If the granule passes TLSCF investigative QA, the SciQAFlag is set to Good, the QA Explanation field is updated using the Metadata Update Tool and explanatory text is provided for the QA Notebook, in the same way as was done for failed granules (Steps 5 and 6).

5.3 Manual QA at the TLSCF - Routine

Even if no granules are marked for investigative QA by the automatic QA process, some percentage of the granules in the production stream will undergo manual QA (Step 10). This review serves to provide an additional check on the production process and is called TLSCF routine QA. The percentage of routinely-sampled granules will generally be less than 10% of the total number of scenes processed per day.

The AST has determined relative sampling rates for each data product based on algorithm complexity and the likelihood of problems with the algorithm. Section 7 contains information about data stream sampling. If necessary, these sampling rates will be updated after launch when more experience is gained in performing QA at the TLSCF.

Based on the AST sampling guidelines, the TLSCF QA Engineer submits queries to EOSDIS approximately daily and orders those granules meeting the query criteria. Typical search parameters might include latitude/longitude of the target, day and/or time of acquisition, and day and/or time of processing. These granules are then evaluated using the standard review procedures that are summarized in Appendix B. TLSCF QA personnel may also develop product-specific procedures in conjunction with the algorithm developers. The results of routine QA at the TLSCF are entered in the QANB.

Delivery to the TLSCF will be made via ftp or via digital media. Media shipment schedules are TBD, with daily or weekly shipment most probable.

6.0 Trend Analysis

Trend analysis is the process of comparing daily granules and metadata to similar datasets from earlier in the mission. This type of analysis is used primarily for troubleshooting during manual QA, but may also be used to track behavior of the AM-1 platform, the ASTER instrument, or a PGE throughout the course of the mission.

Trend analyses by DAAC and TLSCF personnel will rely on reviews of the summary statistics, the Alert Log, and the QANB (discussed in detail in Sections 5.1 through 5.3, and in Appendix A.) Summary statistics also provide key indicators of instrument and algorithm performance, even if no alerts have been triggered. The QANB provides a sequential summary of QA activities at the TLSCF performed in response to the alerts and summary statistics.

TLSCF personnel may also use other data products which are indicators of ASTER performance. For example, periodic noise will be more apparent in the Decorrelation Stretch product than in any other, and random noise will tend to be amplified. This data product can thus be used as a sensitive monitor, over time, of the noise characteristics of the instrument. Other ASTER data and metadata may also prove to be useful in augmenting planned QA products.

7.0 Data Volume and QA Workforce

This section provides estimates of the number of granules associated with each component of TLSCF QA and the workforce required to complete the QA tasks for those granules. The numbers provided are generally "minimum" numbers. For example, value could be added by increasing the sampling density for the Data Stream Sampling component. Also, it is assumed that the Level 1 data stream has already been assessed for quality, and that major Level 1 problems been identified and corrected by Japanese processing teams.

7.1 Data Stream Sampling

Data stream sampling will concentrate on the more complex algorithms and those algorithms that are dependent on external data sources, as these are more likely to have problems. These products are listed in the following table along with the expected daily production rate (based on ECS Ad Hoc Working Group estimates), the number of scenes proposed to be manually assessed (about 10 % of the total) and a rough estimate of the time required to review each scene (this does not include corrective actions). These time estimates assume a strictly visual/cognitive approach to evaluation is adequate (i.e., no software or other evaluation assistance tools are used. In reality, it is likely that some software tools will be employed -- to examine metadata, for example.) These estimates also assume that data production is in "steady-state", i.e., most of the problems associated with the beginning of data production have been resolved.

Because Surface Reflectance is calculated simultaneously with VNIR/SWIR Surface Radiance, the quality of these two products is highly correlated. Thus, manual evaluation will focus on only one, Surface Reflectance.

Product	Granules Generated / Day	Granules Reviewed / Day	Minutes per Scene	Total Review Time
VNIR/SWIR Reflectance	62	7	3	21 min
TIR Radiance	62	7	3	21
Surface Temperature	62	7	3	21
Emissivity	62	7	3	21
Total				84 min

The simpler algorithms and those without external dependencies are expected to be much more reliable and to rarely produce granules of poor quality (unless the 1B input was poor). It is proposed that these granules be reviewed less frequently -- perhaps 1% of the total produced, as outlined in the following table.

Product	Granules Generated / Day	Granules Reviewed / Day	Minutes per Scene	Total Review Time
D-stretch (VNIR)	341	4	3	12 min
D-stretch (SWIR)	341	4	3	12
D-stretch (TIR)	341	4	3	12
Brightness Temperature	62	1	3	3
Total				39 min

Overall, manual evaluation of both groups results in about 2 hours per day required for data stream sampling. This will allow time for further action to be taken for bad scenes, as well as for other QA activities.

7.2 Handling Bad Granules

To estimate the workforce required to handle bad granules, very rough guesses were made as to the number of problems that might occur in a day, and how long it would take to identify and resolve them. This was done for three phases of the mission. Although there will probably be a gradual decrease in the number of problems throughout the mission life, most of the overall decrease is expected to occur in the first nine months.

Launch+1 month to Launch+9 months: This phase is characterized by low data volumes, sometimes one scene or less per day, and usually less than 10, particularly in the first 105 days after launch. On the other hand, the problem rate and the magnitude of the problems are expected to be greater than in the other phases. Because this is the first time the algorithms are being used with real data it is expected that the algorithm developers will play a very active role in assessing quality and correcting problems. In many ways this phase is a mixture of development, validation, and operations, and the TLSCF QA engineer will largely play the role of facilitator. The TLSCF QA engineer's main task may be to connect identified problems with the appropriate developer, and to handle the operational components of the QA process, while the developers handle the algorithmic components. A workforce estimate of 1.0 TLSCF QA person has been made for this phase, with additional support from the algorithm developers and DAAC QA personnel, as needed.

Launch+9 months to Launch+15 months: During this phase it is expected that many problems will still be surfacing and require attention. At the same time, user demand is increasing and there is a gradual increase in the number of granules generated. Due to these factors, it is estimated that there will be two to three times as many problems as in other phases. The TLSCF will maintain 1.0 QA person during this phase, with additional help from algorithm developers and DAAC QA personnel as needed for problem identification and correction.

Plateau Phase: This is the term informally used for the nominal part of the mission after most of the initial problems have been resolved. Data granule throughput will be at its maximum. A very rough estimate is that there will be perhaps one problem per day for each of the three complex algorithms (TIR Radiance, VNIR / SWIR Radiance and Reflectance, and Temperature/Emissivity Separation). Because this problem may be manifested in one or perhaps many granules, resolving it requires both finding and addressing the cause, as well as taking steps to document or re-generate the affected granules. A workforce estimate of 1.0 TLSCF QA person has been made for this phase.

APPENDIX A: Product-by-Product Alert Information

This appendix describes the alert information collected for each ASTER data product on a product-by-product basis. The first six sections contain alert details provided by the algorithm developers about their products. Tables 1 through 3 at the end of this section provide a compilation of the summary statistic and alert information for ASTER Level 2 (L2) products.

Brightness Temperature at Sensor (BTS)

There is a single alert condition to be logged by the BTS algorithm. If pixels have been detected with temperatures outside certain threshold values (tentatively set at -100 deg. C. and +100 deg. C.), this fact is noted. Scenes that trigger this alert need to undergo manual QA to determine whether the values seem appropriate within their scene context. If they seem appropriate (such as a cold pixel in Antarctica or a hot pixel within a forest fire), no further action is needed. If a scene has inappropriately hot or cold pixels, the scene should be forwarded to the algorithm developer for review.

A note about the BTS look-up table (LUT) of temperatures: For the BTS data product, a table is created where brightness temperature is represented as a function of observed radiance. To generate the brightness temperature product for a given scene, one uses the radiance value of each input pixel as the index to point to the desired brightness temperature in the stored table, and place the table values in the output dataset. Alerts occur when temperature values are not within this LUT.

As with all the algorithms, once there is operational experience in evaluating scenes, the threshold values may be adjusted to decrease false alarms.

Decorrelation Stretch (DST)

There are three alert conditions to be logged by the DST algorithm. These alert conditions are:

- 1) an off-diagonal input correlation matrix element is greater than or equal to one
- 2) a computed eigenvalue is less than a threshold value (currently set to 10^{-5})
- 3) an off-diagonal input correlation matrix element is less than a threshold value (currently 0.5)

These alert conditions could occur naturally, but will more likely be due to incorrect or corrupted input data. For example, if two of the three requested input channels are accidentally data from the same channel, an alert would result.

Any DST product that flags an alert condition should be set aside for manual QA. The above threshold values may be adjusted, if it is deemed appropriate.

Atmospheric Correction-TIR (ACT)

The ACT algorithm has one alert currently defined. If additional bad pixels were added by running the ACT algorithm in comparison to the number of bad pixels seen in L1 input products, the granule should be set aside for manual QA. That is, if pixels which were considered good in L1 data products are considered bad in Level 2 data products after running the ACT algorithm, the granule should be set aside for manual QA.

Atmospheric Correction-VNIR/SWIR (ACVS)

There are two summary statistics currently defined for the ACVS algorithm. These are listed below, along with a brief explanation.

- 1) Values in all bands for a given pixel exceed some critical value.

Very few surfaces on the earth exhibit high spectral reflectance in all ASTER bands. If the reflectance in all bands for a given pixel exceeds a given high value, it indicates that the pixel contains specular reflection or clouds and the user should be warned that the results will have high uncertainties in these cases. The critical value will be set initially to a reflectance of 0.8.

The number of pixels from this summary statistic should not exceed 0.1% of the pixels in that scene (a scene is equal to 7.2×10^6 pixels.) If this percentage is exceeded then an alert will be sent.

- 2) The atmospheric correction required extrapolation of the look-up table (LUT).

This indicates that the atmospheric case encountered was outside of the bounds of the look-up table. It is important to note here that the ACVS look-up table provides surface radiances and surface reflectances as a function of top-of-atmosphere radiance. The radiance/reflectance tables are selected based on viewing geometry (sun angle and view angle), and measured atmospheric properties (aerosol type, aerosol optical depth and molecular optical depth).

Temperature Emissivity Separation (TES)

Currently, two alerts are generated by the TES algorithm. The first alert applies specifically to the TES data product. For the case where all TIR bands are utilized, if the algorithm fails for a given percentage of pixels (% To Be Supplied [TBS]), then manual QA is indicated and the algorithm developers wish to be notified. This alert would specifically indicate that there is a problem with the way in which the algorithm is performing T/E separation.

The second alert is the more general case of the percentage of bad pixels for any reason (% TBS). The algorithm developers make a distinction between failure in the

full-band T/E separation and the band-lack T/E separation. The first case indicates a problem with algorithm robustness (provided primarily for the algorithm developers and "power users"), while the second is provided more for user browse information (i.e., is this product worth getting?). This second TES alert can be handled by reviewing the first four bits of the first QA dataplane.

Algorithm testing thus far on simulated ASTER datasets indicates that very few pixels fail to converge, so it is expected that the first alert shouldn't be generated very often. It is recognized that the arrival of actual flight data may alter this assessment. The second alert is expected more often because there are more reasons for which it's generated.

Polar Surface and Cloud Classification (PSCC)

The only PSCC alert identified at this time indicates when there is a large fraction of pixels (value TBS) with low certainty values. This alert would be an indication that a large fraction of the pixel feature vectors was located on decision boundaries. This alert indicates the need for manual QA.

Summary of Product-by-Product Alert Information

The following tables summarize the summary statistic and alert information provided in the previous six sections of this appendix. Table 1 lists the summary statistics and their attributes, Table 2 lists the alerts and their attributes, and Tables 3 through 5 describe alert-related metadata and the QA Summary of Alerts, a text string which provides alert summary information for all the alerts triggered in a particular granule.

Table 1: Summary Statistic Definitions

Algorithm	SumStat Attribute	SumStat Description	Associated Valid Range ¹
Brightness Temperature	QASumStatNum TempImposs	The number of pixels with an impossible temperature	-273.15°C > T > +150°C
	QASumStatNum TempOOR	The number of pixels with a possible but out-of-range temperature	-100°C > T > +100°C
Atmospheric Correction: TIR	QASumStatNum AddlBadPixels	The additional number of bad pixels added by L2 processing	Not Applicable
Atmospheric Correction: VNIR/SWIR	QASumStatNum PixAllBndsOOR	The number of pixels where all bands exceed their critical values	Reflectance ≤ 0.8 or Radiance ≤ Threshold Radiance [†]

¹ The valid range values given here are "at launch" values and are subject to change with experience with the coded algorithms and with the receipt of mission data. These valid range values are stored in a look-up table associated with the algorithm code for ease of access and ease of maintenance; the range values can be easily changed without an algorithm code re-delivery.

†

Band 1: 474.3 X cosine(solar zenith) [W/m²/sr/micrometer]

Band 2: 397.1 X cosine(solar zenith) [W/m²/sr/micrometer]

Band 3: 282.2 X cosine(solar zenith) [W/m²/sr/micrometer]

Band 4: 57.7 X cosine(solar zenith) [W/m²/sr/micrometer]

Band 5: 22.0 X cosine(solar zenith) [W/m²/sr/micrometer]

Band 6: 20.9 X cosine(solar zenith) [W/m²/sr/micrometer]

Band 7: 19.0 X cosine(solar zenith) [W/m²/sr/micrometer]

Band 8: 16.9 X cosine(solar zenith) [W/m²/sr/micrometer]

Band 9: 15.2 X cosine(solar zenith) [W/m²/sr/micrometer]

	QASatLUT ExtrpReqd	The atmospheric correction required extrapolation of the LUT	FALSE
Temperature Emissivity Separation	QASatPctFail Pixels	The % of pixels for which the algorithm fails	Not Applicable
	QASatPctBad Pixels	The % of pixels that are bad in the output product	Not Applicable
Polar Surface and Cloud Classification	QASatPctHi UncertPixels	The % of pixels with high uncertainty values	Not Applicable

Table 2: Alert Definitions

Algorithm	Alert Attribute	Alert Description	Crit?	Trigger Threshold
Brightness Temperature	QAAlertNumTemp Imposs	Summary Statistic exceeded threshold value and triggered this alert	Yes	QAStatNum TempImposs ≥ 1 pixel
	QAAlertNumTemp OOR	Summary Statistic exceeded threshold value and triggered this alert	No	QAStatNum TempOOR > 50 pixels
Decorrelation Stretch	QAAlertOffDiag ElemOOR	Any off-diagonal input correlation matrix element is greater than or equal to one	Yes	Not Applicable
	QAAlertEigenval OOR	Any of the three eigenvalues is less than a threshold value	No	Eigenval $< 10^{-5}$
	QAAlertOffDiagInp LTLim	An off-diagonal input correlation matrix element is less than a threshold value	No	Element < 0.5
Atmospheric Correction-TIR	QAAlertNum AddlBadPixels	Summary Statistic exceeded threshold value and triggered this alert	Yes	QAStatNum AddlBad Pixels ≥ 1 pixel
Atmospheric Correction-VNIR,SWIR	QAAlertNumPixAll BndsOOR	The number of pixels where all bands exceed their critical values exceeds threshold	Yes	QAStatNumPix AllBndsOOR $> 0.1\%$ of the pixels in a scene (where a scene = 7.2×10^6 pixels)
	QAAlertLUTExtrpReqd	The atmospheric correction required extrapolation of the LUT	Yes	Trigger if TRUE
Temp. Emiss. Separation	QAAlertPctFail Pixels	Summary Statistic exceeded threshold value and triggered this alert	Yes	QAStatPctFail Pixels $> TBS$
	QAAlertPctBad Pixels	Summary Statistic exceeded threshold value and triggered this alert	Yes	QAStatPctBad Pixels $> TBS$

Polar Surface and Cloud Classification	QAAlertPctHi UncertPixels	Summary Statistic exceeded threshold value and triggered this alert	Yes	QASummaryofAlerts UncertPixels > TBS
--	------------------------------	---	-----	---

Table 3: Alert-Related Metadata

Attribute Name	Description
QACritAlertsCnt	Number of critical alerts for this granule.
QANonCritAlertsCnt	Number of non-critical alerts for this granule.
QASummaryofAlerts	A text string containing summary information on all the alerts triggered for this granule. Format of this attribute can be found in Table 4.

Table 4: Format of QASummaryofAlerts

Name	Description	Critical?	Actual Value	Valid Range
QAAlertxxxxxxxx	User-friendly text description	Yes / No	Numerical value returned with alert.	y.y > Val > z.z
Must be less than 30 characters.	320 characters	3 characters	10 characters	25 characters

N.B. The format of the QA Summary of Alerts table will be similar to that shown in Tables 4 and 5, though exact specifications (e.g., column width) may vary slightly when coding is completed.

Table 5: Example of QASummaryofAlerts

Name	Description	Critical?	Actual Value	Valid Range
QACritAlertsCnt	Number of critical alerts for this granule.	No	1	Not Applicable
QANonCritAlertsCnt	Number of non-critical alerts for this granule.	No	1	Not Applicable
QAAAlertNumAddlBadPixel	L2 processing added bad pixels	Yes	1034	≤ 100
QAAAlertEigenvalOOR	Eigenvalue is less than threshold value	No	10^{-6}	$\geq 10^{-5}$

APPENDIX B: Product-by-Product Manual QA

The following sections provide information about manual QA for each ASTER data product, as provided by the algorithm developers.

For All Products

For each of the ASTER data products listed below, visual inspection will be performed, providing a simple "sanity check" in order to detect gross errors in the images. Experience with previous sensors indicates that the following errors are likely to be encountered:

- 1) "ringing", where, as you approach a point on the image, values periodically gradually get brighter then plunge to the darkest possible values.
- 2) gradients, where, as you move across (or down) an image, the scene more or less uniformly grows darker (or brighter), without any indication that the target surface should behave in this manner.
- 3) checkerboarding, where, at regular intervals, usually one to a few pixels, the scene alternates from darker than average to brighter than average.
- 4) speckles, where individual pixels seem to have anomalous values, somewhat like a paint splatter.

In general, the QA analyst should be looking for anything that appears to be an artifact of the processing, rather than an honest reporting of the nature of the target surface. This is necessarily a matter of judgment, with some experience and/or practice required.

Brightness Temperature (BTS)

The Brightness Temperature product relies solely on visual inspection, as described above.

Decorrelation Stretch (DST)

While visual inspection of the Decorrelation Stretch products may detect gross errors in the images, it is unlikely to detect most of the possible types of errors from the Decorrelation Stretch program itself, since these errors are usually apparent in the image statistics, rather than in the images themselves. For this reason, a statistical scan of all images set aside for manual QA will be performed. A simple software routine can compute the mean and variance of each of the three output image planes, and the correlation matrix among the three inputs and three outputs combined (a 6x6 matrix), then compare this with the anticipated results.

Atmospheric Correction -TIR (ACT) and Atmospheric Correction - VNIR / SWIR (ACVS)

The manual QA for the atmospheric correction products will provide an additional test of reasonableness for the image outputs. The generating program provides a certain amount of checking on a pixel-by-pixel basis, but often the context within a scene provides a much more stringent test for reasonable values.

Typically, the output image for each channel will look nearly identical to the corresponding input image for that channel. The output image is expected to have a somewhat higher level of contrast, but features within the scene should retain their relative brightness (darker objects remain darker, brighter objects remain brighter, no ringing), should retain their location and clarity (no movement or blurring), and should be free of the introduction of new features (no artifacts of processing).

The downwelling radiance images corresponding to each channel will appear to be very similar for each of the five images. The extreme channels (10 and 14) should be somewhat darker than the central channels, but otherwise similar in appearance. Each image will appear to be an inversion of the corresponding digital elevation model image (i.e., higher elevations appear darker), perhaps with an overlaying brightness trend or gradient. There should be no discontinuities in the pixel values throughout the image.

Atmospheric Correction VNIR/SWIR (ACVS) Only

This section provides additional product-specific manual QA criteria for the ACVS algorithm. By examining specific sets of flagged pixels, certain conclusions may be drawn about data quality.

If all of the flagged pixels with reflectances below the critical value are in the middle of a lake then the developers know they have a problem with the characterization of the aerosols (probably overestimated).

If there are a lot of flagged pixels that are too bright in several bands all in the same place, there is a likely cloud contamination problem or there is a specular reflection off a water surface which mimics the cloud problem.

A case where the look-up table (LUT) bounds are exceeded indicates that the algorithm misused the inputs, the inputs are bad, or the algorithm generated the LUT incorrectly. (The ACVS section of Appendix A provides additional information about the contents of the look-up table.)

Finally, random scatter of the flagged pixels with little correlation between bands would indicate the algorithm missed the spectral characteristics of the aerosols.

Temperature Emissivity Separation (TES)

Reviewing the first 16 bits of the emissivity metadata in the second QA dataplane will reveal why the algorithm failed; that is, if one or more emissivities were out of range, or which land-leaving radiance or sky irradiance input values were bad or suspect. The first four bits of the first QA dataplane would provide additional information whether or not the problem was TES-specific or of a more general nature.

If regular alerts of a general "bad" nature are generated, and manual QA determines this is because Level 1 data are routinely "bad" or because there are too few good bands to perform T/E separation, no further action should be taken. On the other hand, if the "TES algorithm failure" alert triggered manual QA, the algorithm developers would like to know the results of the metadata review.

Polar Surface and Cloud Classification (PSCC)

Manual QA of the PSCC products will consist of comparing a color coded version of the polar cloud mask with either gray level displays of individual bands or three-band color overlays of selected bands (e.g., Band 13=Red, Band 4=Green, and Band 1=Blue). Based on this visual comparison, QA personnel can determine how well the classification mask matches up with what was seen in the raw image data. QA personnel are assumed to have some knowledge of what clouds, land, and other classifications look like in each band of the raw imagery and which three-band overlays are the best to use for PSCC QA.

Though this QA method is simple, classification problems will usually be obvious. If a classification problem is observed, the only way to correct it is to notify the algorithm developer that later re-training of the algorithm is needed to handle the problematic pixels.

QA personnel should also be aware that the algorithm developer does not consider any combination of reflectances and temperatures to be unreasonable, based on the Landsat TM data that are being used for algorithm development and testing. For instance, reflectance values greater than 1.0 are not physically possible, but due to uncertainty about surface aspect and slope, values greater than 1.0 can be obtained, especially on bright snow covered surfaces.

APPENDIX C: ASTER QA Data Planes

ASTER First QA Data Plane

Table 1 summarizes the Good, Suspect and Bad categories and includes annotation for each category. In using this table, please note the following:

- Categories annotated as "All Level 2 (L2) Algorithms" will apply to all higher-level data products (with the possible exception of DEMs).
- Additional categories applicable to a specific algorithm (e.g., Temperature Emissivity Separation (TES) algorithm only) will be used only for that algorithm.

The first column in the table describes the quality category: Good, Bad, or Suspect. The second column provides the binary code for the quality category. Because there are 4 bits available, there are 16 binary codes available. Finally, the third column of the table provides the text of the Good/Bad/Suspect descriptions which have been provided by the ASTER Science Team.

Tables 2 and 3 describe the cloudiness and cloud adjacency pixel codes, respectively, and also map binary codes to particular pixel descriptions.

Table 1: ASTER First QA Data Plane

Note: While the first QA data plane is intended to be generic for all products, the bad and suspect categories do contain product-specific information. This was done because space was available and it deemed a reasonable compromise to accommodate the needs of these products.

Category	Binary Code	Description
Good	0000	This pixel has no known defects.
Suspect	0001	All L2 algorithms: Perimeter effect from thin cloud.
Suspect	0010	All L2 algorithms: Perimeter effect from thick cloud.
Suspect	0011	TES only: Some TES output bands are out-of-range.
Suspect	0100	DEM only: Edited DEM pixel.

Suspect	0101	All L2 algorithms: Algorithm or LUT returned "suspect input value" flag.
Suspect	0110	All L2 algorithms: Output Data value is Out-of-Range.
Suspect	0111	All bands of the input pixel are suspect (reason not specified).
Bad	1000	Saturation: At least one band is saturated.
Bad	1001	Skew: This is a skew (i.e., border) pixel.
Bad	1010	TES only: Too few good bands; no values for T or e are reported.
Bad	1011	TES only: Algorithm divergence. NEM T and e are reported.
Bad	1100	TES only: Algorithm convergence failure. Normalized Emissivity Method [NEM] T and e are reported.
Bad	1101	All L2 algorithms: Algorithm or LUT returns "bad input value" flag
Bad	1110	All L2 algorithms: General LUT failure.
Bad	1111	All L2 algorithms: This pixel was Bad in the L1B product.

Table 2: Cloudiness pixel codes in the First QA Data Plane FOR VNIR ONLY

Binary Code	Description
00	Clear
01	Thin Cloud
10	Thick Cloud
11	Not Used

Table 3: Cloud Adjacency pixel codes in the First QA Data Plane

FOR VNIR ONLY

Binary Code	Description (see Note)
00	Far (Less than or equal to 12 pixels in an 11x11 window contain clouds)
01	Slightly Near (More than 12 and less than or equal to 30 pixels in an 11x11 window contain clouds)
10	Near (More than 30 and less than or equal to 36 pixels in an 11x11 window contain clouds)
11	Very Near (More than 36 pixels in an 11x11 window contain clouds)

Note: To calculate cloud adjacency, a filter window 13 pixels by 13 pixels in size is moved across the data one pixel at a time. Cloud adjacency is calculated based on the number of cloudy pixels within the window each time the window is moved. For example, if 55 pixels in the 11x11 window contain clouds, roughly half of the 121 pixels in the window are cloudy and the center pixel is labelled Very Near a cloud. This cloud adjacency calculation is an intrinsic part of the creation of the first QA data plane and is thus performed for each granule. Values in Table 3 are subject to change when actual flight data is received and processed.

ASTER Second QA Data Plane

The second QA data plane contains product-specific QA information on a pixel-by-pixel basis. The size of the second QA data plane is 8 bits or 16 bits. The following sections provide the details of the second QA data plane for each ASTER data product. Please note that some products do not have a second QA data plane.

In all the second QA data plane tables, LSB denotes the Least Significant Bit and MSB denotes the Most Significant Bit.

Brightness Temperature Separation

The Brightness Temperature Separation algorithm does not have a second QA data plane.

Decorrelation Stretch

The Decorrelation Stretch algorithm does not have a second QA data plane.

Atmospheric Correction TIR

The Atmospheric Correction TIR algorithm assigns 16 bits in the second QA data plane according to Table 1.

Table 1: Assignment of bits in second QA plane for Atmospheric Correction TIR

Emission Channels Uncertainty										Channel Bad or Suspect					Not Used
LSB															MSB
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Channel 10		Channel 11		Channel 12		Channel 13		Channel 14		Ch 10	Ch 11	Ch 12	Ch 13	Ch 14	Not Used

The two bits for the emission uncertainty for each channel will follow the following format:

Binary Code	Meaning
00	Uncertainty < 2%
01	Uncertainty 2-5%
10	Uncertainty 5-15%
11	Uncertainty > 15%

Atmospheric Correction VNIR/SWIR

The Atmospheric Correction VNIR/SWIR algorithm assigns 16 bits in the second QA data plane according to Table 2.

Table 2: Assignment of bits in second QA plane for Atmospheric Correction VNIR/SWIR

Reflection Channels Uncertainty											Channel Bad or Suspect				
LSB											MSB				
0	1	2	3	4	5	6	7	8	9	10	11	12	13	14	15
Channel 4	Channel 5	Channel 6	Channel 7	Channel 8	Channel 9	Bit Assignments TBS by Algorithm Developer									

The two bits for the reflection uncertainty for each channel will follow the following format:

Binary Code	Meaning
00	Uncertainty < 5%
01	Uncertainty 5-10%
10	Uncertainty 10-20%
11	Uncertainty > 20%

Temperature Emissivity Separation

The Temperature Emissivity Separation (TES) algorithm produces both a temperature and an emissivity product, as the name suggests. The first eight bits of the second QA data plane are used to provide QA information common to both the temperature and the emissivity products, as shown in Tables 3 and 4.

It is important to note that the values in Table 4 are correct as of this writing and are subject to change when actual flight data is received and processed. These values are stored in the product-specific metadata for each product.

Table 3: First 8 Bits Common to both Temperature and Emissivity

0	1	2	3	4	5	6	7
LSB						MSB	
Maximum Emissivity, binned (ϵ_{max})		Number of iterations		Sky Irradiance/Land-Leaving Radiance ratio ($S \downarrow / L'$)		Minimum emissivity reset flag	Not Used

Table 4: Description of QA Bits Common to both Temperature and Emissivity

Parameter	Values	Binary Code	Implication
ϵ_{max}	≤ 0.94	00	possible error condition
	$> 0.94-0.96$	01	most silicate rocks
	$> 0.96-0.98$	10	default value of ϵ_{max}
	> 0.98	11	vegetation, snow, water, some soils
Number of Iterations	4	00	fast to converge
	5	01	nominal performance
	6	10	nominal performance
	≥ 7	11	slow to converge
$S \downarrow / L'$	≤ 0.1	00	high altitude scene; correction probably accurate
	$> 0.1-0.2$	01	nominal value
	$> 0.2-0.3$	10	nominal value
	> 0.3	11	warm, humid air or cold land; correction may be inaccurate
ϵ_{min} Reset Flag	Not reset	0	ϵ_{min} did not need to be corrected
	Reset	1	ϵ_{min} was reduced in proportion to measurement noise

The remaining bits provide information about either the Temperature or Emissivity data product and are appended only to the appropriate product (Tables 5 and 7). The Temperature and Emissivity bit descriptions are shown in Tables 6 and 8, respectively.

Table 5: Assignment of bits in second QA plane for Temperature Data Product

0	1	2	3	4	5	6	7
---	---	---	---	---	---	---	---

LSB		MSB
Accuracy	Precision	Band Used for Calculating Temperature

While bits are available for Accuracy and Precision for Temperature products, these values are not set by the processing software and zero-fill is used for these bits. Accuracy can only be calculated by comparing ASTER data with ground truth data, and therefore bit descriptions are not provided in this document.

Table 6: Temperature Product QA Bit Descriptions

Binary Code	Band Used for Calculating Temperature
0000	ASTER Band 10 used (not normally used)
0001	ASTER Band 11 used
0010	ASTER Band 12 used
0100	ASTER Band 13 used
1000	ASTER Band 14 used

Table 7: Assignment of bits in second QA plane for Emissivity Data Product

8	9	10	11	12	13	14	15	16	17	18	19	20	21	22	23-31
Ch 10		Ch 11		Ch 12		Ch 13		Ch 14		Not Used					

In addition to the 8 bits common to both the Temperature and Emissivity data products, the Emissivity data product uses an additional 24 bits to describe errors detected in emissivity calculations. Bit assignments and descriptions are provided in Tables 7 and 8.

As noted above for the Temperature data product, bits are available for Accuracy and Precision for Temperature products, though these values are not set by the processing software and zero-fill is used. Accuracy information will be supplied in bits 24 and 25 and Precision information will be supplied in bits 26 and 27, when these parameters are implemented.

Table 8: Emissivity Product QA Bit Descriptions

Binary Code	Errors Detected
000	Good from good input
001	Good from suspect input

100	Bad or cloudy input
101	Not good from good input
110	Not good from suspect input
111	Intentionally lacked (for band-lack algorithm)

Polar Surface and Cloud Classification

The Polar Surface and Cloud Classification product does not have a second QA data plane.