CHAPTER 3

UNIT FUNCTIONAL FLOW

3.1 <u>Introduction</u>. The purpose of this chapter is to describe the flow of data through the application software of the major functional areas of the WSR-88D unit to enable the user to better understand and manage the available data resources. Emphasis is placed on the flow in the RPG.

3.2 <u>Radar Data Acquisition</u>. The RDASC program is used to control the real-time operation of components of the RDA. The RDASC program monitors and assesses the performance of the RDA, initiates automatic calibration, performs calibration calculations, and reports RDA status to the RPG functional area. It formats reflectivity, mean radial velocity, and (velocity) spectrum width data input from the signal processing subsystem, attaches header data, and initiates data transfer to the RPG. The RDASC program is partitioned into the following functional capabilities (Figure 3-1):

- Handle Maintenance Console,
- Monitor and Calibrate RDA,
- Form RDA State,
- Handle Wideband,
- Control RDA,
- Control Signal Processor, and
- Control Pedestal.

3.2.1 Handle Maintenance Console. This function provides the RDASC Program interface with the Maintenance Console Keyboard and the Data Acquisition Unit (DAU). Inputs to the Handle Maintenance Console function are received from the Keyboard, Monitor, Calibrate RDA function, the DAU, and the Control RDA function.

Command entries from the keyboard are validated to ensure that they conform to the allowable set in format and type. Associated parameter entries are validated to determine if they are within the allowable range. When the RDA is in the remote control state, certain commands can only be issued from the RPG. Invalid entries result in error messages being displayed; valid entries are acknowledged on the display. Command and parameter entries are made through the use of menus that prompt the user for inputs.

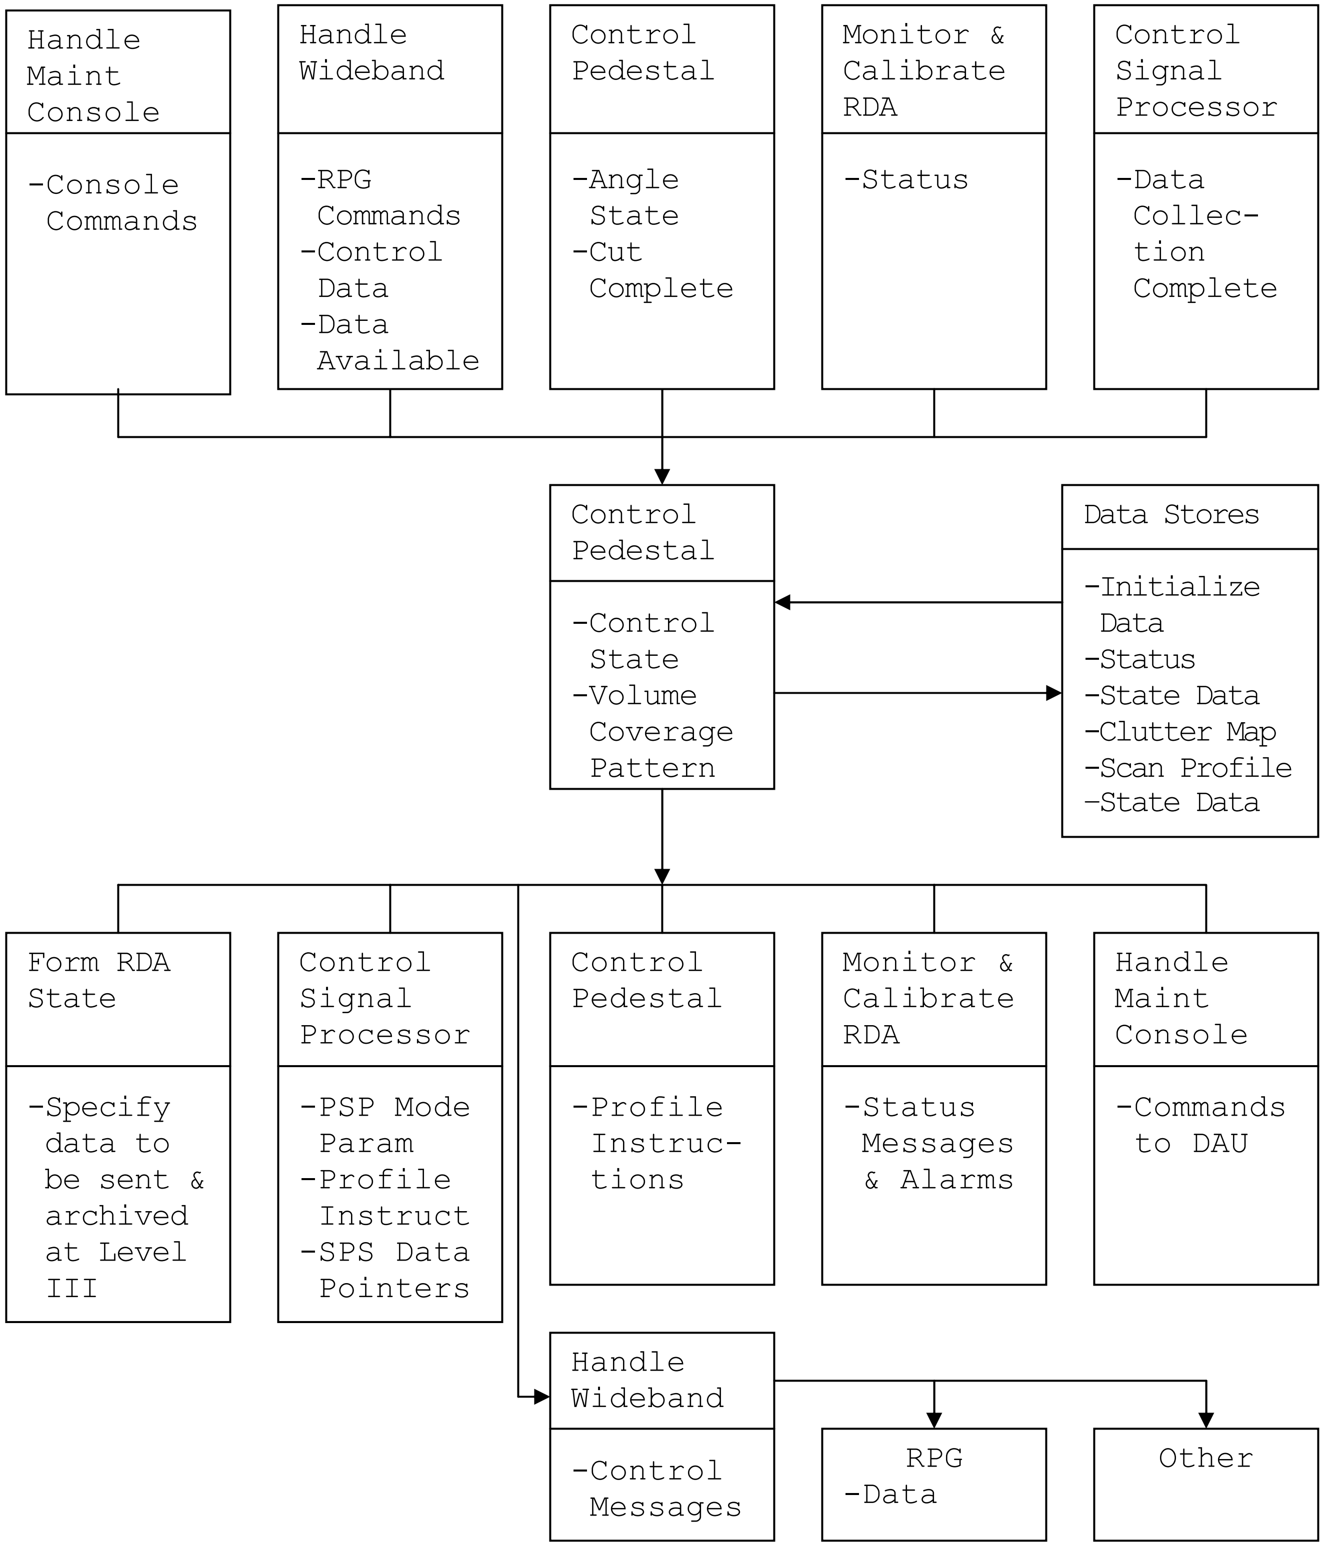

Figure 3-1
Radar Data Acquisition Status and Control Program Data Flow

An Enter Log Data command provides for the entry of data into a maintenance log file. The entry procedure includes prompts that assure that complete information is input for each event to be logged. A Display Log Data command provides for the display of previously entered data. The capacity of the log is 100 records with an overload warning issued when the log contains 90 or more.

An Inspect/Change RDA Adaptation Data command allows the user to inspect or change adaptation data. Changes to certain data are only allowed under password protection with two levels of protection provided. (A listing and description of site-adaptable parameters applicable to the RDASC are contained in and specified by *WSR-88D Operator Handbook, Guidance on Adaptable Parameters, Volume 3, RDA*. The handbook can be obtained at: http://www.roc.noaa.gov/ssb/sysdoc/Operations.asp. The authority to change these parameters is addressed in Part A of this Handbook.) When adaptation data are changed, the original version of the data is retained with changes used to generate a current version. Both versions are maintained on disk and the capability to edit either version is provided.

Summary RDA status and alarm messages, as designated by the Monitor and Calibrate RDA function, are displayed in dedicated and protected areas of the alphanumeric screen. These displays are active at the same time that menus and other data are displayed. The status display is updated automatically when changes in status occur. Performance data can be displayed upon command with separate menus provided for the various categories of data.

When a Cold Startup or Restart command is issued from the Control RDA function, actions are taken automatically to initialize the DAU and Maintenance Console function. When the DAU and Maintenance Console interfaces are initialized, the DAU initialization status and the Maintenance Console initialization status contained in the RDA performance data are set to OK or Fail, depending on whether or not the initialization was successful. If an error occurs during the initialization process, appropriate alarms are set.

Outputs from this function are directed to the Maintenance Console and its alphanumeric display, to the RDA Data function, to the Data Store, and to the Control RDA function.

3.2.2 Monitor and Calibrate Radar Data Acquisition. This function generates RDA performance data based on Built-In Test (BIT) data and other inputs from the RDA equipment, other program functions, and the Operating System. Inputs consist of pedestal angle state information from Control Pedestal function and signal processor, transmitter, receiver, rf generator BIT data, and tower and utilities data from the Data Store. Data Store also provides scan information for scheduling tests and calibration data from the signal processor.

The function maintains the RDA summary status at the Maintenance Console for transmission to the RPG through the Form RDA Data function, whenever changes in the data occur and when RDA status is requested by the RPG. The RDA performance data are also formed at the

Maintenance Console for transmission to the RPG as they are generated by this and other functions of the RDASC program.

3.2.2.1 Evaluate Hardware Status. The RDA performance data are generated for each hardware configuration item based on inputs from the DAU, pedestal, and receiver and signal processor. The DAU BIT inputs are available from the Handle Maintenance Console function. Pedestal and receiver and signal processor data are received from the Control Pedestal and Control Signal Processor functions. Calibration data are received from the Perform RDA Calibration subfunction. RDA alarms are set when discrete BIT data indicate a fault or when analog BIT data or calibration data are out of tolerance.

3.2.2.2 File Maintenance. Whenever data are read from a disk file or written to a disk file, the disk file read/write status is updated based on the status reported by the Operating System; it is then included in the RDA performance data. Disk files are maintained for the RDA state file, bypass map, calibration data, adaptation data, censor zone, and remote VCP.

3.2.2.3 Perform Radar Data Acquisition Calibration. Calculations of calibration parameters are performed when RDA calibration processing is scheduled by the Control RDA function. This activity generates applicable RDA performance data, performs calibration checks, and generates alarms, when appropriate. It is performed during the period of elevation transition between volume scans or elevation cuts and at startup.

Outputs from the Monitor and Calibrate RDA function (summary status, performance data, alarms, and calibration parameters) are sent to the Handle Maintenance Console and Control Signal Processor functions.

3.2.3 Form Radar Data Acquisition Data. This function assembles data to be output on the wideband communications links, inserts required header data, generates the associated data parameter blocks, and passes the data parameter block to the Control Wideband function for data transmission. Data are input from the Control Signal Processor, Monitor and Calibrate RDA, Control RDA, and Control Pedestal functions.

When notified of data to be output on the wideband data links, the Form RDA Data function generates a header for the data, which includes message size, message type, an identification sequence number, a date and time tag, the number of separate message segments that make up the data message, and the individual segment number. Other header information for base data includes pedestal angle positions and a calibration parameter that is derived from automatic calibration routines or RPG input. Outputs include a parameter block for each data set to be transmitted, header data segments to be incorporated into the data sets, and notification to the Handle Wideband function that data are available.

3.2.4 Handle Wideband Data. This function supports and controls the wideband communications link to the RPG. Inputs include RPG to RDA data received over the wideband link, header data, and an RDA data pointer list from the Form RDA Data function, message data from data stores, wideband state from the wideband communications handler, and wideband control from the Control RDA function. This function is divided into three subfunctions: Handle Wideband Link, Transmit Wideband Data, and Receive Wideband Data.

3.2.4.1 Handle Wideband Link. This subfunction provides an interface with the channel terminal manager, which is part of the Operating System. It provides the capabilities to establish and terminate communication links and to control transfer of data between the link interfaces and other functions of the RDASC program.

3.2.4.2 Transmit Wideband Data. When notified by the Form RDA Data function that radials of data are ready to be transmitted, this subfunction determines whether sufficient buffers are available for the Control Signal Processor Subsystem (SPS) function to continue reading data out of the PSP. If sufficient buffers are not available, it determines which interface (RPG, wideband user, or Archive II) is responsible for the unavailability of sufficient buffers and deactivates it to prevent it from interfering with data transmission over the other interfaces.

3.2.4.3 Receive Wideband Data. When notified by the Handle Wideband Link subfunction that a message has been received, this subfunction notifies the Control RDA function that data are available.

The Handle Wideband function outputs include RPG to RDA data received over the wideband link, RDA to RPG data and RDA to user data to be transmitted over the wideband link, wideband status data from the wideband communications handler, and wideband control from the Control RDA function.

3.2.5 Control Radar Data Acquisition. The Control RDA function performs processing in response to control commands entered at the RDA Maintenance Console or received from the RPG through the Handle Wideband function. It also provides overall sequence control of the volume coverage by initiation of processing within the Control Signal Processor and Control Pedestal functions based on the selected VCP. Inputs include entries from the RDA Maintenance Console and the MSCF, hardware status information from the hardware interfaces or other RDASC functions, and control data from the RPG. This function is further divided into two subfunctions:

3.2.5.1 Control Radar Data Acquisition State. The RDA states are combinations of control states and operating states. Control states include local (RDA), remote (RPG), and either (RDA or RPG). Operating states include standby, restart, startup, playback, offline operate, and operate. Control states define the commands that can be accepted from the RDA Maintenance Console or RPG while operating states define the processing that is performed by the RDA in response to commands. Transitions between the various control states and operating states can take place due to commands from the RDA or RPG, or automatically when specified conditions occur.

3.2.5.2 Volume Coverage Control. The Volume Coverage Control subfunction initiates and controls the execution of the selected VCP as well as the RDA calibration functions. The selected VCP is determined by the command issued at the RDA Maintenance Console or the RPG. If no command was issued, the default VCP identified in the RDA state file is selected. If a new VCP is selected while a volume scan sequence is in progress, the new VCP takes effect when the next volume scan starts.

Outputs of the Control RDA State include status messages and alarms to the Handle Maintenance function; status data, commands, and parameters to other RDASC program functions; and transmitter control commands.

3.2.6 Control Signal Processor. This function loads the SPS controls into the SPS and receives and unpacks the SPS outputs. The SPS input data consist of base data and surveillance power data, calibration data, on-line performance monitoring data, signal processor loads data, and control RDA function inputs.

When power is first applied to the SPS at startup or restart following a power loss, microcode required by the PSP is automatically downloaded to the PSP by the SPS. This download may also be commanded from the RDASC processor.

3.2.6.1 Clutter Filter Data Transfer. During normal radial processing, the proper radial of the Clutter Suppression Map is downloaded to the SPS. The elevation dimension of the map data is taken from the elevation angle for the cut in the VCP. The azimuth dimension of the map data is the pedestal angle projected forward to account for pipeline delays. The Clutter Suppression Map is derived at startup and whenever a new set of Clutter Suppression Regions are received by the RDA (downloaded by the MSCF). This map is derived by merging the default suppression levels defined by the Default Notch Width Map and clutter locations contained in the Bypass Map definition with the operator-defined Clutter Suppression Regions. When compiling the Clutter Suppression Map, operator-defined Clutter Suppression Regions take precedence for the application of clutter suppression; however, the Bypass Map and Default Notch Width Map control the application of suppression for every range bin NOT contained within an operator-defined region. The actual notch width values assigned to each radial of the Clutter Suppression Map are determined by the notch width verses scan rate and suppression level table maintained in RDA adaptation data, where the scan rate is determined by the VCP and the suppression level is determined from the level defined in the Clutter Suppression Map.

3.2.6.2 Compute Interference Detection Rate. When the Interference Detection Unit is configured, the function computes the interference detection rate contained in the Summary RDA Status.

Outputs to the RDASC resident portion of the Signal Processing Program consist of macrocode setup parameters. These parameters are transferred either by shared memory or in the form of subroutine parameters. Receiver and signal processor control and supplemental data are downloaded to the SPS through Direct Memory Access transfer.

Output from Data Stores includes on-line performance monitoring data from the SPS, receiver, and RF generator, calibration data used by the Monitor and Calibrate RDA function, and digital radar data that are made available to the Handle Wideband function.

Other outputs are digital radar data buffer pointer consisting of address information needed to access digital radar data that are sent to the Form RDA Data function; a calibration parameter used by the PSP to calculate reflectivity from signal power, which goes to the Form RDA Data function; notification to the Control RDA function indicating that the SPS data report has been received and the status of the SPS; and a flag used to initiate transition to the next elevation angle scan, which is sent to the Volume Coverage Control subfunction.

3.2.7 Control Pedestal. The Control Pedestal function of the RDASC program provides three basic and essentially separate subfunctions. The primary function is to perform closed loop servo control of the pedestal azimuth and elevation positions. This control may be by either rate or position. Scan Profile Data input to the Control Pedestal function provides the basic information describing the desired volume scan scenario. The Control RDA function instructs the Control Pedestal function as to which set of data is to be used for a particular volume scan. It also provides real-time commands to commence or terminate a volume scan or, when appropriate, to restart the scan at any particular point after a short-term power outage. In return, Control Pedestal informs Control RDA when particular angle or angle rate transition sequences are stabilized so that radar data processing can be initiated.

A secondary function of Control Pedestal is to provide azimuth and elevation angular positions to other functions. These angles are appended to Digital Radar Data by the Form RDA Data function. The remaining function of Control Pedestal is to receive pedestal built in test equipment (BITE) data from the pedestal and forward these data to the Monitor and Calibrate RDA function.

Inputs include scan profile data that define the information needed to perform various volume scan scenarios; pedestal data, which include pedestal azimuth and elevation angular positions; and angular rates as well as pedestal BITE data. Profile instruction input is provided from the Control RDA function as well as pedestal position correction parameters.

When a startup command is issued by the Control RDA function, pedestal self tests are conducted by issuing appropriate commands to the pedestal and checking the test results.

3.2.7.1 Generate Angular Position. This subfunction generates the azimuth and elevation position commands for the pedestal as a function of time. The scan patterns are specified in terms of constant rate or position segments, or both, separated by transition events. Operational scan strategies specified to date require only constant elevation continuous azimuth rotation patterns. For this type of pattern, the segments are constant elevation positions with a constant rate azimuth rotation and the transition event is the start of a new cut.

3.2.7.2 Control Pedestal Servo. The position of the pedestal in azimuth and elevation is controlled by this subfunction. It compares present pedestal position to commanded pedestal position and generates rate commands to the pedestal based upon position error and desired rate.

3.2.7.3 Handle Pedestal Input and Output. Communications with the pedestal in support of all pedestal control functions is provided by the Handle Pedestal Input and Output. These communications are implemented over a serial data link through a controller tied to the RDASC computer.

3.2.7.4 Format Pedestal Built-In Test Equipment. The Format Pedestal BITE continually updates pedestal state data, which can be accessed by the Monitor and Calibrate RDA function for purposes of ascertaining the operability state of the pedestal.

Outputs of the Control Pedestal function include pedestal control commands, consisting of commands regarding the azimuth and elevation drive rates and pedestal data requests sent to the pedestal through the Handle Pedestal Input/Output (I/O); angle state, consisting of data and flags that are sent to the Control RDA function; and pedestal state, consisting of pedestal BITE data and a fault flag that are accessed by the Monitor and Calibrate RDA function.

3.3 Radar Product Generation. The RPG functional area controls the operation of the entire WSR-88D unit. It is organized into modules that, taken together, accomplish the following functions:

- System Status and Control,
- Base Data Acquisition,
- Preprocess Base Data,
- Product Generation, and
- Product Distribution.

Figure 3-2 contains a summary of the flow and functionality of each of these modules. The purpose of these modules is to transform digitized weather data received from the RDA, supplemental external data, and user input data into base products, derived products, alphanumeric products, and data arrays in forms suitable for archive and distribute to WSR-88D users.

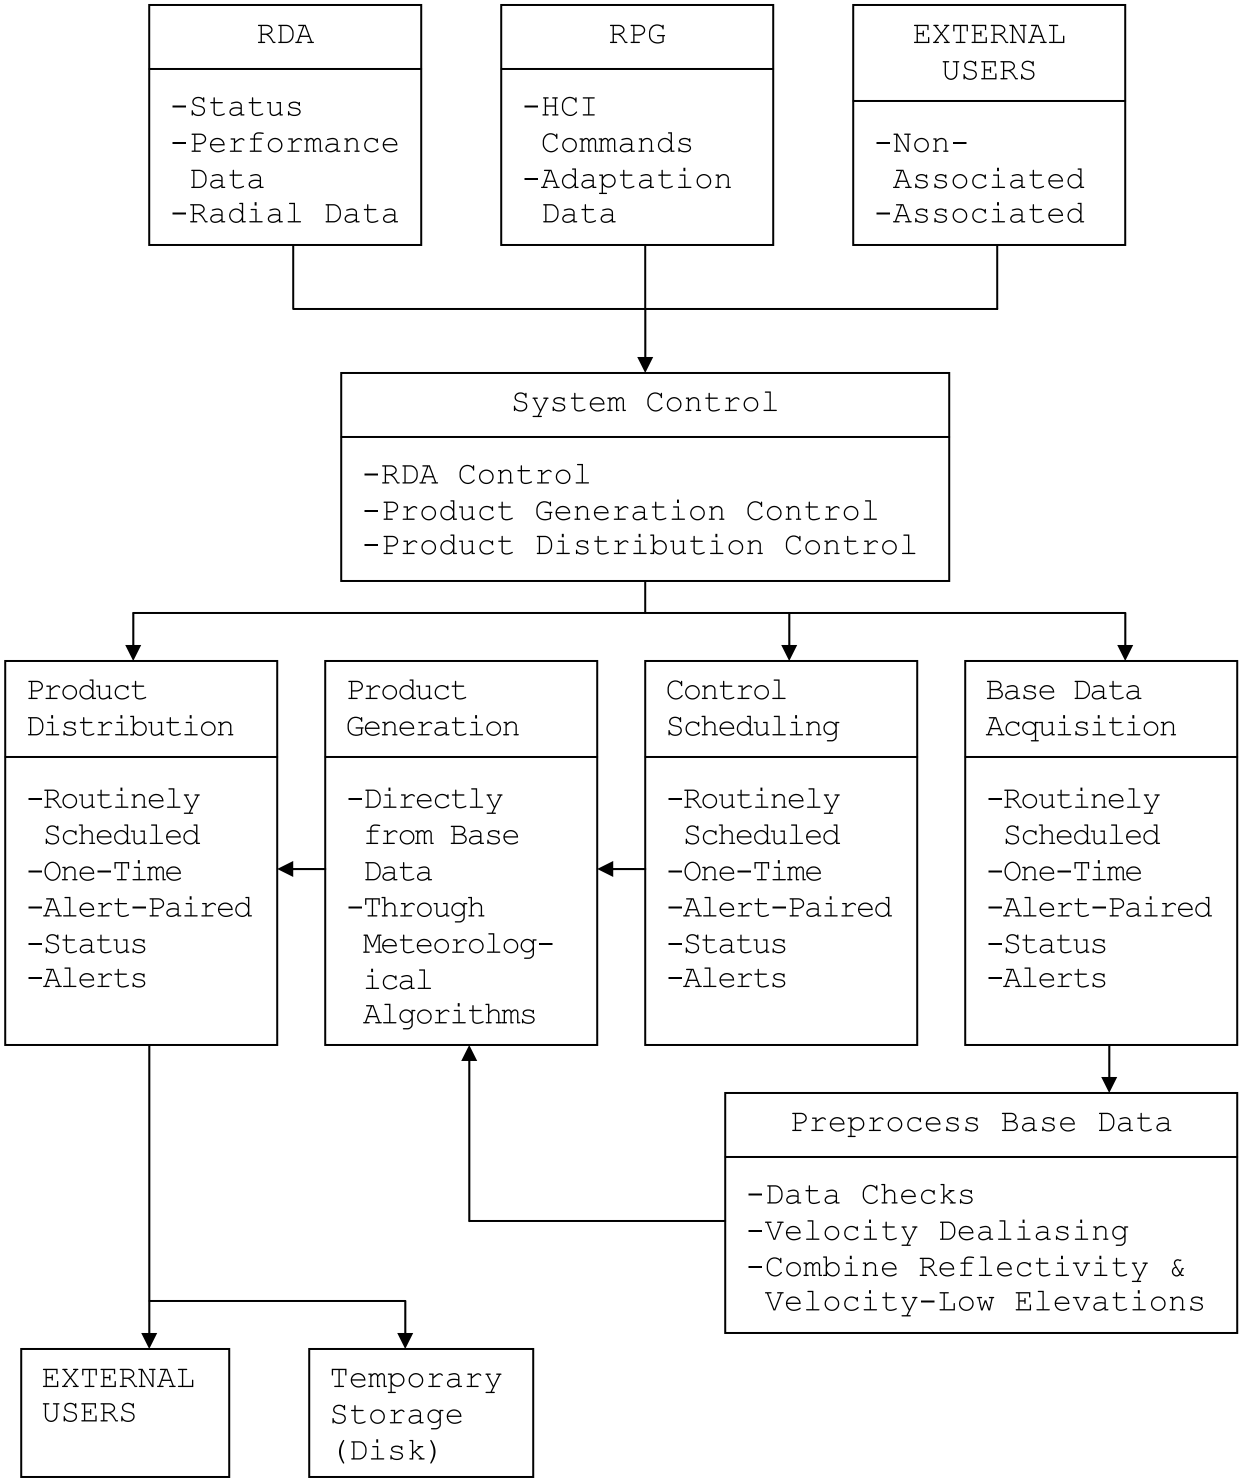

Figure 3-2
Radar Product Generation Data Flow

3.3.1 System Control. Within the RPG, the System Status and Control function accepts input data from the RDA, the RPG itself, and display systems. Site-adaptable parameters resident within the RPG are listed in the WSR-88D Operator Handbook, Guidance on Adaptable Parameters. As mentioned earlier, authority for change of these parameters is addressed in Part A of this Handbook. These data are used for both RDA control, RPG control, and product generation and distribution control.

The System Status and Control function also presents to the system user information concerning overall system health through a combination of RDA and RPG status displays and RDA Performance Data.

3.3.1.1 System Status and RDA Control. The RPG user controls the RDA through the use of RDA Control Commands available at the RPG Human Computer Interface (HCI). Operationally, the most frequently used commands are those specifying VCP, PRFs used within the VCP, and Clutter Suppression Region definitions.

Other commands available allow the system user to restart the VCP, change between utility and backup power source, place the RDA in Standby or Operate State, restart RDA application software, and change the RDA control state between Local RDA Control and Remote RDA Control. The RPG HCI user can only control the RDA when the RDA is in Remote RDA control.

The RPG HCI displays RDA status information to enable the HCI user to monitor the RDA for alarm messages and performance quality. Data available to the user includes transmitter, receiver/signal processor, antenna/pedestal, and communications performance.

3.3.1.2 Product Generation. At the beginning of each volume scan, the RPG creates a table of products that will be produced for that volume scan. This table is made up of products specified at the MSCF on the Generation and Distribution Control Menu, products in response to alert thresholds being exceeded, Routine Product Set (RPS) lists input from Associated User display systems, and one-time requests from Associated User display systems.

3.3.1.3 Product Distribution. Distribution of products is controlled according to inputs from User display systems and the MSCF as described in Section 3.3.5.

3.3.1.4 Product Generation Control. Generation of products is controlled according to inputs from External User Systems and the RPG HCI as described in Section 3.3.4.

3.3.1.5 Product Distribution Control. Distribution of products is controlled according to inputs from External User Systems and the RPG HCI as described in Section 3.3.5.

3.3.2 Base Data Acquisition. The Base Data Acquisition function is responsible for managing the wideband communications link to the RDA and processing incoming and outgoing RDA/RPG interface control messages. Messages passed from the RDA to RPG and processed by the Base Data Acquisition function include Digital Radar Data, RDA Status Data, Performance/Maintenance Data, Console Message, Clutter Filter Bypass Map, and Clutter Filter Notchwidth Map. Messages passed from the RPG to the RDA and processed by the Base Data Acquisition function include RDA Control Commands, Volume Coverage Pattern, Clutter Censor Zone Data, and Console Message.

3.3.2.1 RDA and RPG Radial Loadshedding. In the RPG, there are two queues where radial data is stored and monitored. In one of the queues, radial data in RDA/RPG ICD format (i.e., Archive II/Level II format) is buffered. The other queue contains the radial data after conversion from interface control format to internal RPG format (this latter format is for RPG algorithm processing convenience). Thus, the two queues represent different locations within the radial processing stream of the RPG. Figure 3-3 contains a summary of RPG loadshedding.

The current utilization values presented for these categories are the ratio of number of unread (i.e., unprocessed) radials in each queue to the maximum number of radials each queue can hold. Unlike the other loadshed categories, if the Alarm Levels are exceeded, no actual loadshedding mechanism is invoked unless queue overflow occurs. A loadshed Alarm merely indicates the processing of radial messages is severely backlogged and product generation based on radial data input may fail. Product generation failures will occur if either of the queues overflows because newly arriving radial messages will replace unread radial messages in the queues.

3.3.3 Preprocess Base Data. The RPG performs certain preprocessing functions on the base data before making these data available for meteorological algorithm processing and product preparation. This preprocessing includes: performing angle calculations, ensuring angle changes are within tolerance, velocity dealiasing, and other continuity considerations. The processed radials are then released to downstream consumers. In addition, these processed radials are saved in a database so that one-time requests for products not routinely produced can be handled within a reasonable time period.

3.3.3.1 Velocity Dealiasing. Velocity dealiasing is performed using two methods, the first uses an algorithm (Part C of this Handbook) in the RPG in order to take advantage of the processing power there and facilitate the implementation of improved dealiasing algorithms in the future. The MPDA is also executed at this time in the RPG. We will begin with the original method. The following description is intended only to present the main functional elements of the complex algorithm.

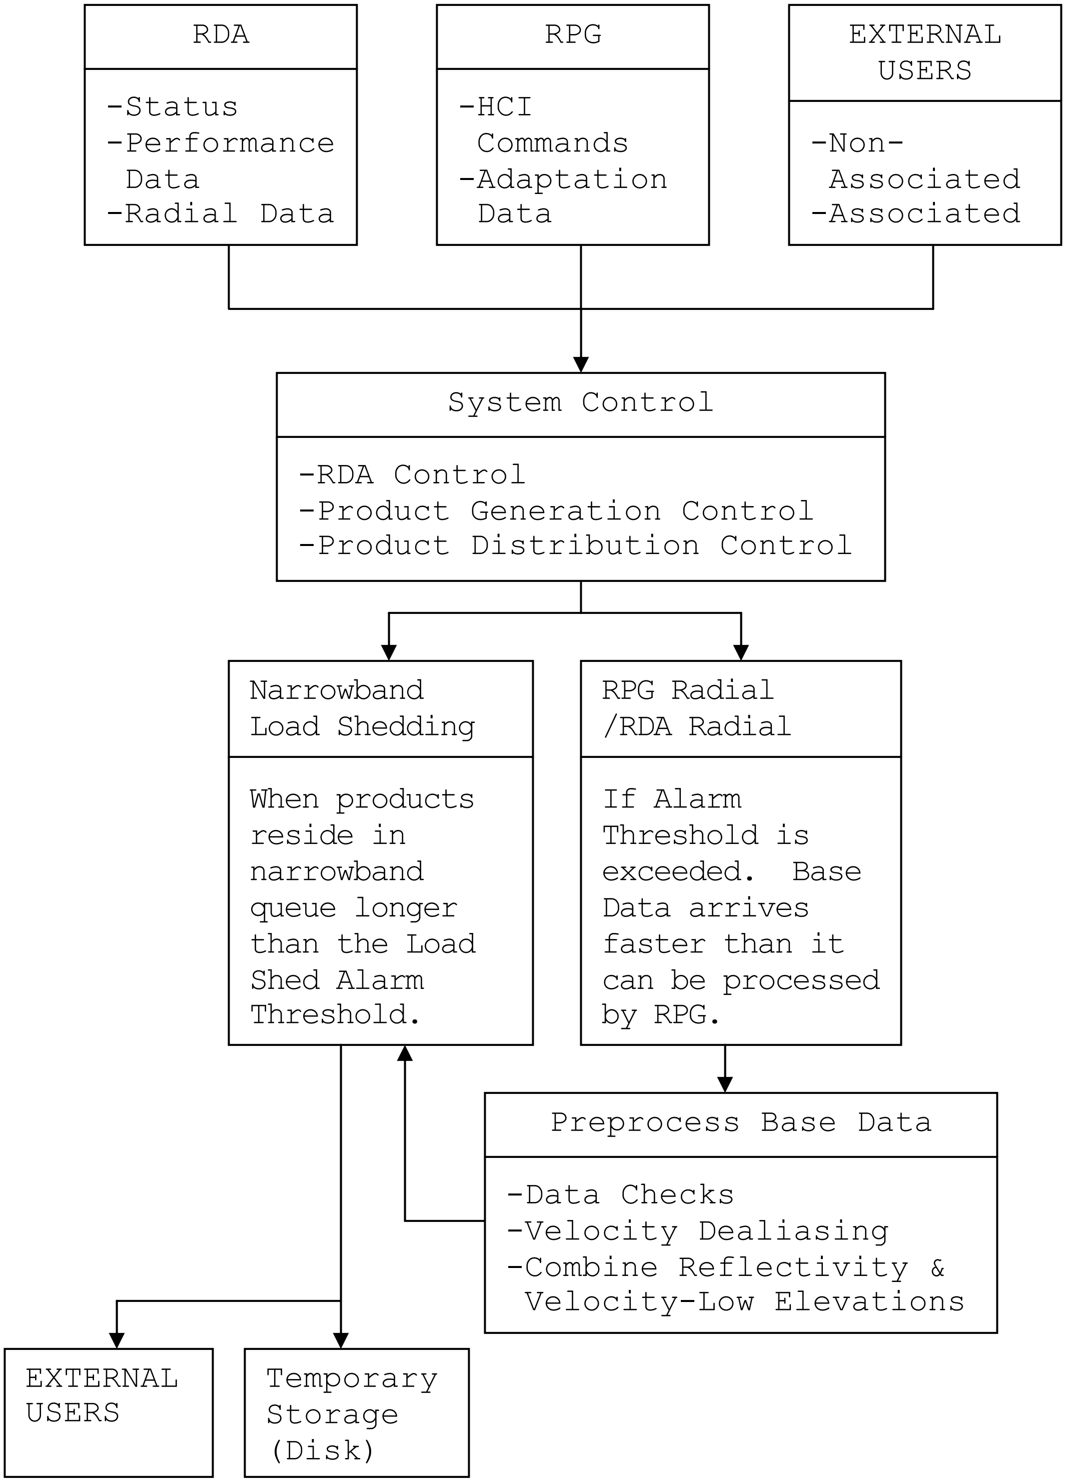

Figure 3-3
Types of Load Shedding at the Radar Product Generator

The algorithm dealiases one radial of velocity data at a time, operating from the first range bin out to 124 nm. Each velocity value is checked for spatial continuity both with previous bins, which have been dealiased on the same radial, and with already dealiased bins near the same range on the previous radial. The algorithm allows for large changes in velocity values if the changes occur over limited distances. This provides the ability to avoid falsely dealiasing the velocity values associated with meteorological features such as Mesocyclones and Tornado Vortex Signatures. If spatial continuity checks indicate an unacceptable velocity value and the value is spatially isolated, the algorithm uses environmental wind information (entered at the HCI and updated by the Velocity Azimuth Display Algorithm) to attempt to dealias the suspect value.

If neither spatial continuity nor environmental wind checks enable the dealiasing of the suspect value, the value is rejected. Rejected velocity values are temporarily stored for possible reinsertion into the radial. If five velocity values in a row are rejected, they will be reinserted into the radial. If a radial cannot be dealiased completely, it is not used for azimuthal continuity checks on the next radial's velocity values. In this case, the previous completely dealiased radial continues to be used for the azimuthal continuity checks. If five successive radials cannot be dealiased completely, no further azimuthal continuity checks are made during the rest of that scan's dealiasing process. This feature prevents unsuccessfully dealiased areas from propagating azimuthally around the scan.

The second method uses the MPDA approach. The MPDA dealiasing scheme is a multi-step process that arrives at a final dealiased velocity solution at each radar gate. Throughout the processing steps described below, “seed” velocities are used to check the gate-by-gate MPDA results for consistency. These seeds can be single previously dealiased velocity solutions along the same radial, averages of previously dealiased gates, or estimates from the Environmental Wind Table (EWT).

The main dealiasing steps are:

a) Solutions from Velocity Triplets (tight constraint)

This first step considers only gates at which three velocity estimates are present. These estimates for the same gate are obtained by sampling with three different PRFs. The three estimates must be dealiased within a small velocity difference of each other and within a threshold velocity difference of a seed velocity as previously defined for a final solution to be accepted. In general, this processing accounts for about 57% of the final dealiased field.

b) Solutions from Velocity Triplets (relaxed constraint)

The second step considers only gates at which three velocity estimates are present. The three estimates and a seed velocity must be dealiased within a larger velocity difference than step a) of each other.

c) Solutions from Pairs within Triplets

The algorithm next attempts to dealias velocity pairs that did not pass the dealiasing tests in steps a and b. The use of velocity pairs may be due to there being only two estimates at a particular location in space or due to the failure to find a solution using triplets. In the case where triplets are present three solutions are possible. However, the first pair that provides an acceptable solution is retained. Note that at 4.3° elevation, where only two velocity cuts are obtained, MPDA processing begins with this step.

d) Solutions from Single Estimates

At this point the remaining solutions are derived from the single velocity estimates that exist within the unsolved triplets and pairs, and those locations in space that only contained one estimate. The single estimates are dealiased using seed values from the previous steps and increasingly relaxed thresholds. Once this processing is complete, more than 99.99% of the gates contain a final dealiased velocity value.

e) Use original velocity estimates

The remaining gates are assigned one of the three original velocity estimates that are closest to an average of the surrounding dealiased gates.

Error Mitigation Schemes

After steps a through d described above are complete, error mitigation is applied to check for outliers and for azimuthal and radial inconsistencies in the dealiased field.

a) Despeckling

A despeckling function is applied after each processing step. This routine checks for single velocity gates whose solutions differ significantly from surrounding gates. Several averages of surrounding gates are checked against the gate in question. If the gate can be dealiased within a strict threshold of one of the averages, it is assigned a value. Otherwise, it is set to missing and waits further processing.

b) Azimuthal Error Correction

This routine searches for runs of gates along radials that differ significantly when compared to adjacent azimuthal values. Adjacent azimuths on both sides of the azimuth in question are considered in the checking. If the azimuth in question can be dealiased within a strict threshold of its adjacent azimuths, values are assigned to it. Otherwise, the gates on the radial are set to missing and processed further.

c) Radial Error Correction

This routine searches for large gate-to-gate jumps along radials. If a jump is encountered, an attempt is made to dealias it into the correct Nyquist interval based on averages of other radially adjacent gates within a strict threshold. If the gates in question can be dealiased, new values are assigned, otherwise the gates are set to missing and wait further processing.

3.3.3.2 Combining Reflectivity and Doppler Data. To help achieve optimal clutter filter performance, the lowest elevation angles of a VCP (usually 0.5° and 1.5°) are designated split cuts. That is, the RDA collects reflectivity (also called surveillance) data and Doppler data on separate scans. The RPG combines these surveillance and Doppler data into a format consistent with that for higher elevation scans. For a given Doppler data radial, the combining process involves selecting data from the surveillance radial closest in azimuth position to the azimuth position of the Doppler radial.

3.3.4 Product Generation. The RPG Product Generation function coordinates the generation of routine and one-time requested products. Routine requests include products to be generated by default as well as those products requested on a routine basis from External Users via RPS lists. One-time requests (OTRs) include products requested on a one-time basis via OTR messages from External Users, as well as those products generated based on an Alerting criteria being met (Section 3.3.4.2.1). The RPG user can control which products are generated by default by editing the RPG Product Generation Table at the RPG HCI.

3.3.4.1 Routine Product Scheduling. The Routine Product Scheduler is aware of algorithm and product dependencies based on information defined in two configuration tables: the Product Attributes Table (PAT) and the Task Attributes Table (TAT). The PAT lists the types and attributes of data or products the RPG can produce. Attributes include dependent data type(s), the data producer (algorithm), and data-dependent parameters. The TAT lists for each RPG algorithm the data type(s) input to the algorithm and the data type(s) produced by the algorithm.

For each product request, the Routine Product Scheduler uses information contained in the PAT and TAT to determine what data types must be produced in order to generate the requested product. As an example, if Algorithm 1 ingests data type A and produces data type B, and Algorithm 2 ingests data type B and produces product C, then a request for product C requires data types A and B be produced. Hence, a request for product C results in additional requests for data types A and B.

At the beginning of each volume scan, the Routine Product Scheduler builds a Master Request List composed of all algorithm outputs (data types or products) which must be generated in order to satisfy default product requests, External User product requests, unsatisfied one-time product requests, and the Alerting function. After the Master Request List is built, the list is published by the Routine Product Scheduler for all algorithms to read.

The RPG uses a data-driven model for algorithm processing. During the course of a volume scan, every algorithm in the RPG periodically polls whether its input data has been generated for that volume scan. If the input data has been produced, the algorithm then checks the Master Request List for requests for one or more of the algorithm's outputs. If input is available and output is requested, the algorithm processes the input data to produce the output data. Once produced, the output data is written to either an intermediate data store (to be read by down-stream consumers), the product database if the output is a product or the data is to be used to support one-time requests, or in the case of radial data, stored in the Replay Database.

3.3.4.2 One-Time Product Scheduling. When the date and time field is left blank in a one-time request, the Product Generation Function queries the product database for the product in the current volume scan. If a match is found, information about the product is sent to the Product Distribution function (Section 3.3.5). If no match is found, the Product Generation Function queries the product database for the product in the previous volume scan. If a match is found, information about the product is sent to the Product Distribution function. Otherwise, the Replay List is checked to see if the product can be generated on a one-time basis. For those one-time requests that can be generated, the RPG checks whether the product is to be generated from the Replay Database or from warehoused intermediate data. The product is then scheduled for generation through a rerun of base data or warehoused intermediate data. If the one-time requested product can not be generated, the product request is scheduled for generation during the next volume scan as a routine product request.

3.3.4.2.1 Alert Processing. Whenever an Associated External User establishes connection to the RPG, the external system automatically uploads to the RPG the Alert Request message. This message defines the locations and categories for which the External User wishes to be alerted. The External User can modify and resend the Alert Request Message at any time during connection to the RPG.

During the course of a volume scan, the Alert Processing function ingests data generated by RPG algorithms. These data contain the necessary information for alert checking of all alert categories from all Associated External Systems. A check is made for alert conditions for each requested category corresponding to the presently available data. If an alert condition is encountered, an Alert Message is sent to the appropriate external system. If the External User requested that the alert-paired product associated with the alert be sent upon detection of the alert condition, the Alerting Processing function issues a one-time request on behalf of the External User for the alert-paired product. By requirement, only products that are or can be generated using data from the volume scan in which the alert occurred can be paired to an alert.

3.3.4.2.2 Replay Data Products. Product requests (one-time and alert-paired) that are not already generated, and are too late to be processed on the current stream of base (radial) data, may be generated by a replay of radial data or warehoused intermediate data. Due to the enormous amount of processing that would be required, not all volume-based products are included on the replay list. The following product-generating tasks make up the current Replay List (refer to Figure 3-4 (a through g) for task-product relationships):

Mean Radial Velocity	8-bit Velocity
Combined Shear	Radar Coded Message
Reflectivity	8-bit Reflectivity
Composite Reflectivity Grid	Severe Weather Analysis
Spectrum Width	User Selectable Layer Reflectivity
Cross Section	Storm Relative Mean Radial Velocity
Velocity Azimuth Display	

If a one-time requested product is not available and the product is on the Replay List, the product will be produced using replay data. Depending on whether the generating task ingests radial data or intermediate product data determines which database is accessed for the information. Product generating tasks requiring radial data, access the radial replay database. Products requiring intermediate data access the database warehousing intermediate data.

3.3.4.3 Data Flow in RPG. Figure 3-4 (a through g) depicts the flow of data for the generation of products, other than the User Alert Message, the Combined Attributes Table, and the Radar Coded Message, which are composed of the outputs of other products, as indicated in Section 3.3.4.3.3. For a description of the individual outputs, refer to Part C of this Handbook.

3.3.4.3.1 Products Produced Directly from Base Data. A list of products produced directly from preprocessed radial data along with a summary of the data processing follows. These products do not require processing by meteorological algorithms prior to formatting for distribution. One exception is the Severe Weather Analysis product, which receives input from the Shear algorithm to construct the map for radial shear data.

Reflectivity. This task obtains buffers for the six Reflectivity product types (depending on the types requiring generation). Radial data that are input to the task are run-length encoded for the 1 km (0.54 nm) x 1° version. For 2 km (1.1 nm) x 1° versions, the radial is compacted by saving the largest value of two consecutive bins, after which the radial is run-length encoded. For 4 km (2.2 nm) x 1° versions, the radial is compacted twice prior to run-length encoding. The task processes the 1 km (0.54 nm) products first, the 2 km (1.1 nm) next, and the 4 km (2.2 nm) last. These products can be produced for any or all elevation cuts within a VCP.

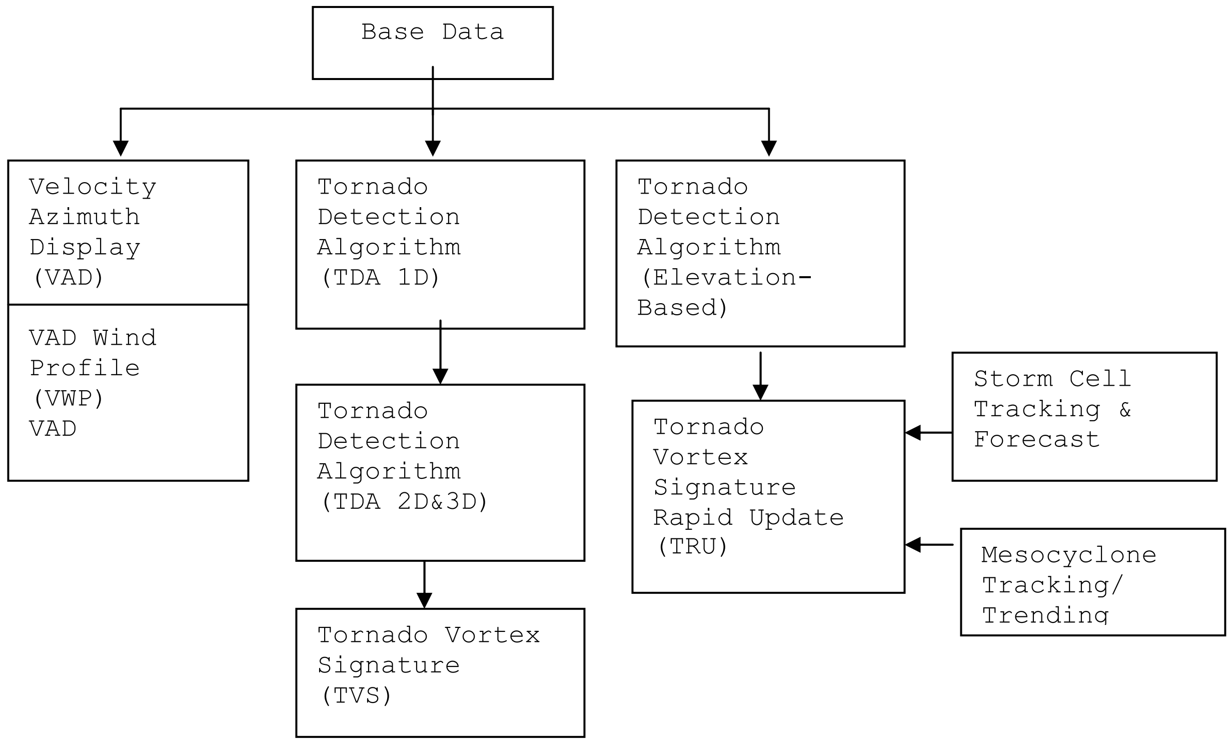

Figure 3-4a
Flow of Data for Generation of Products

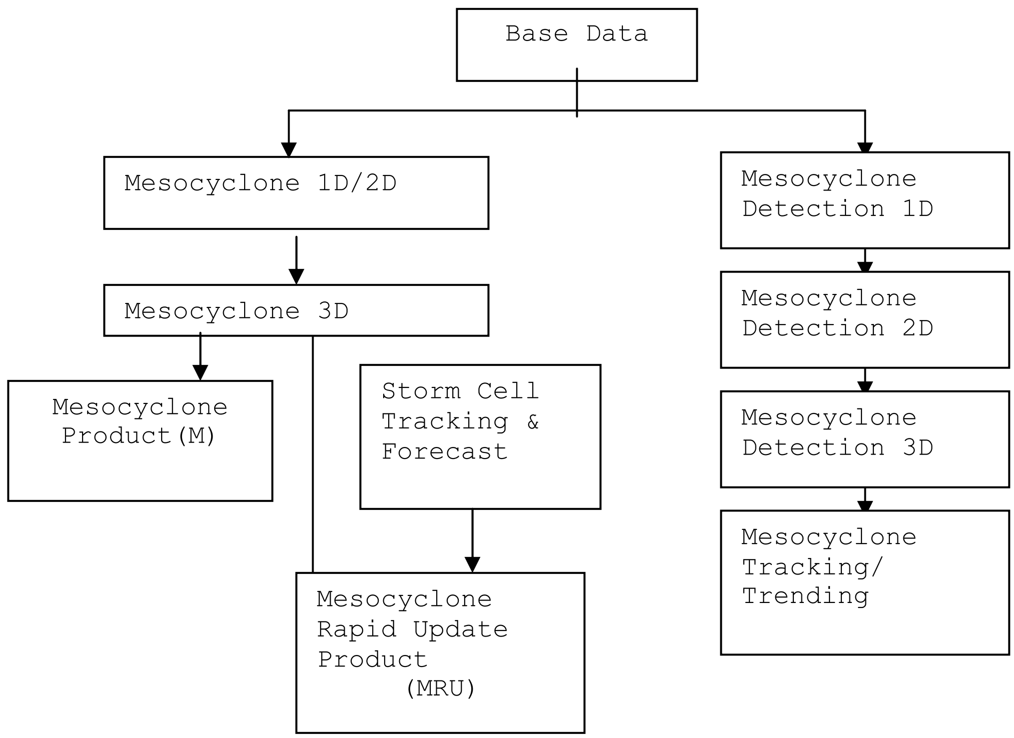

Figure 3-4b
Flow of Data for Generation of Products (Continued)

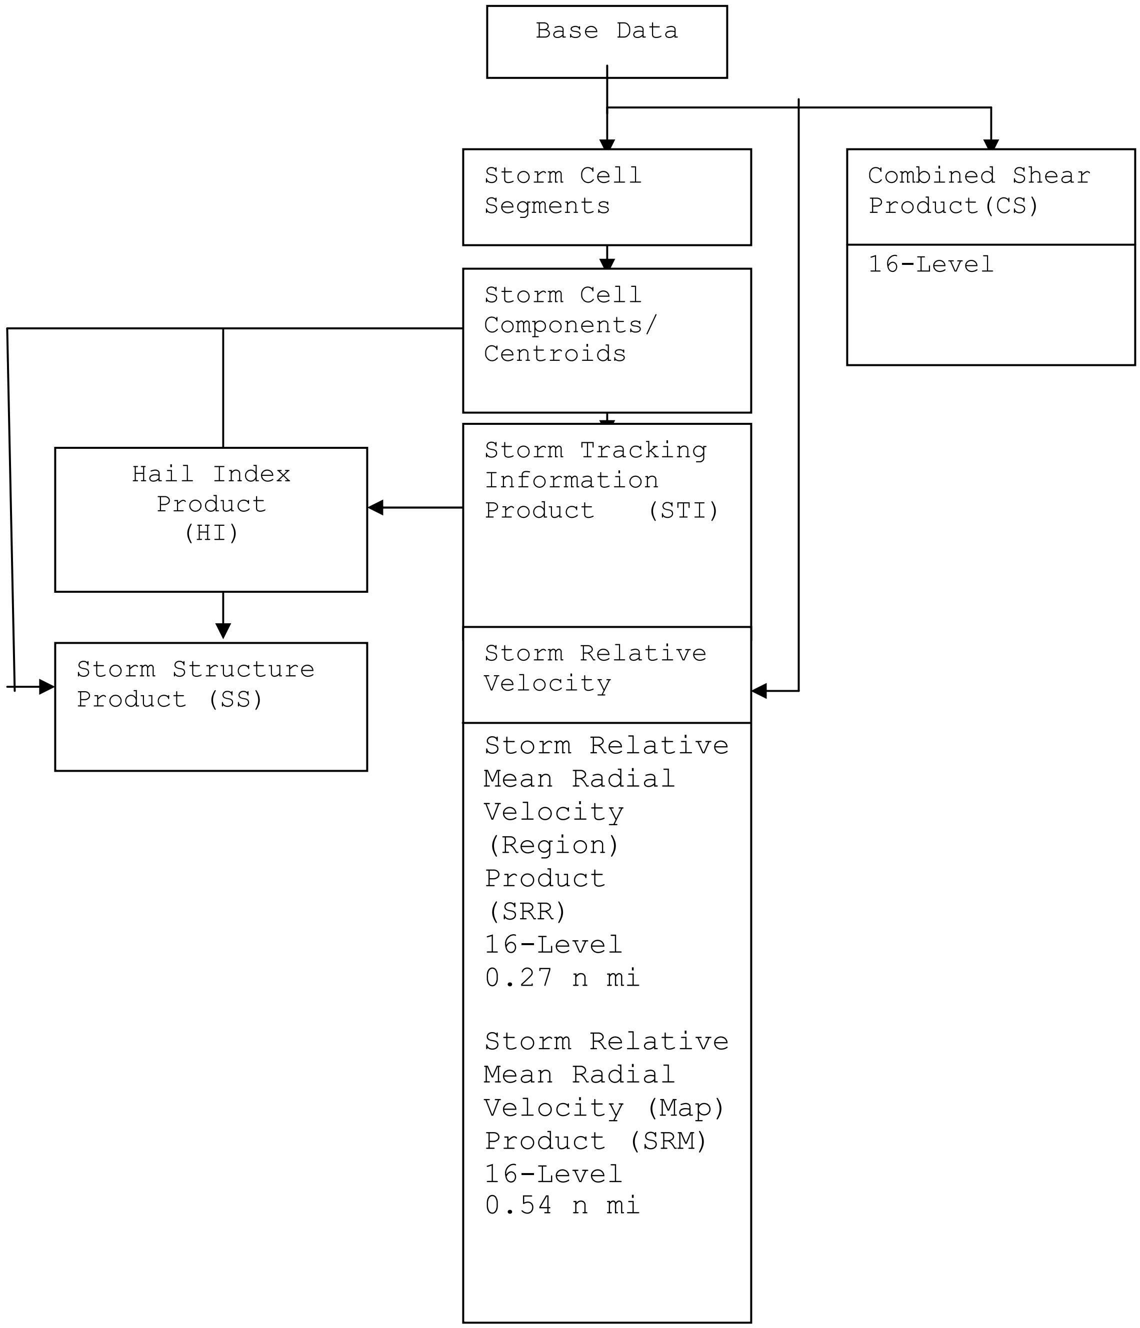

Figure 3-4c
Flow of Data for Generation of Products (Continued)

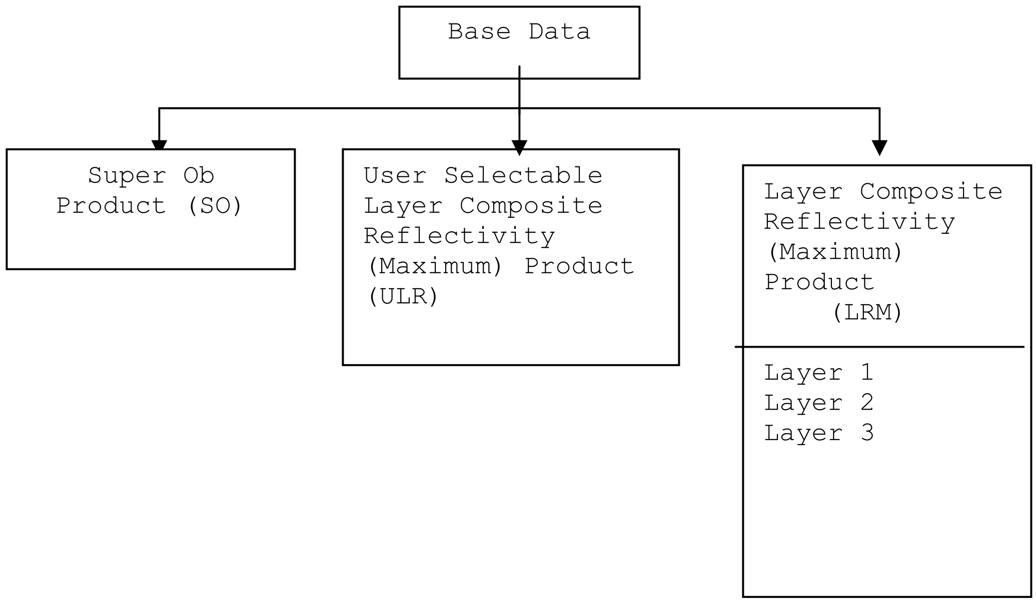

Figure 3-4d
Flow of Data for Generation of Products (Continued)

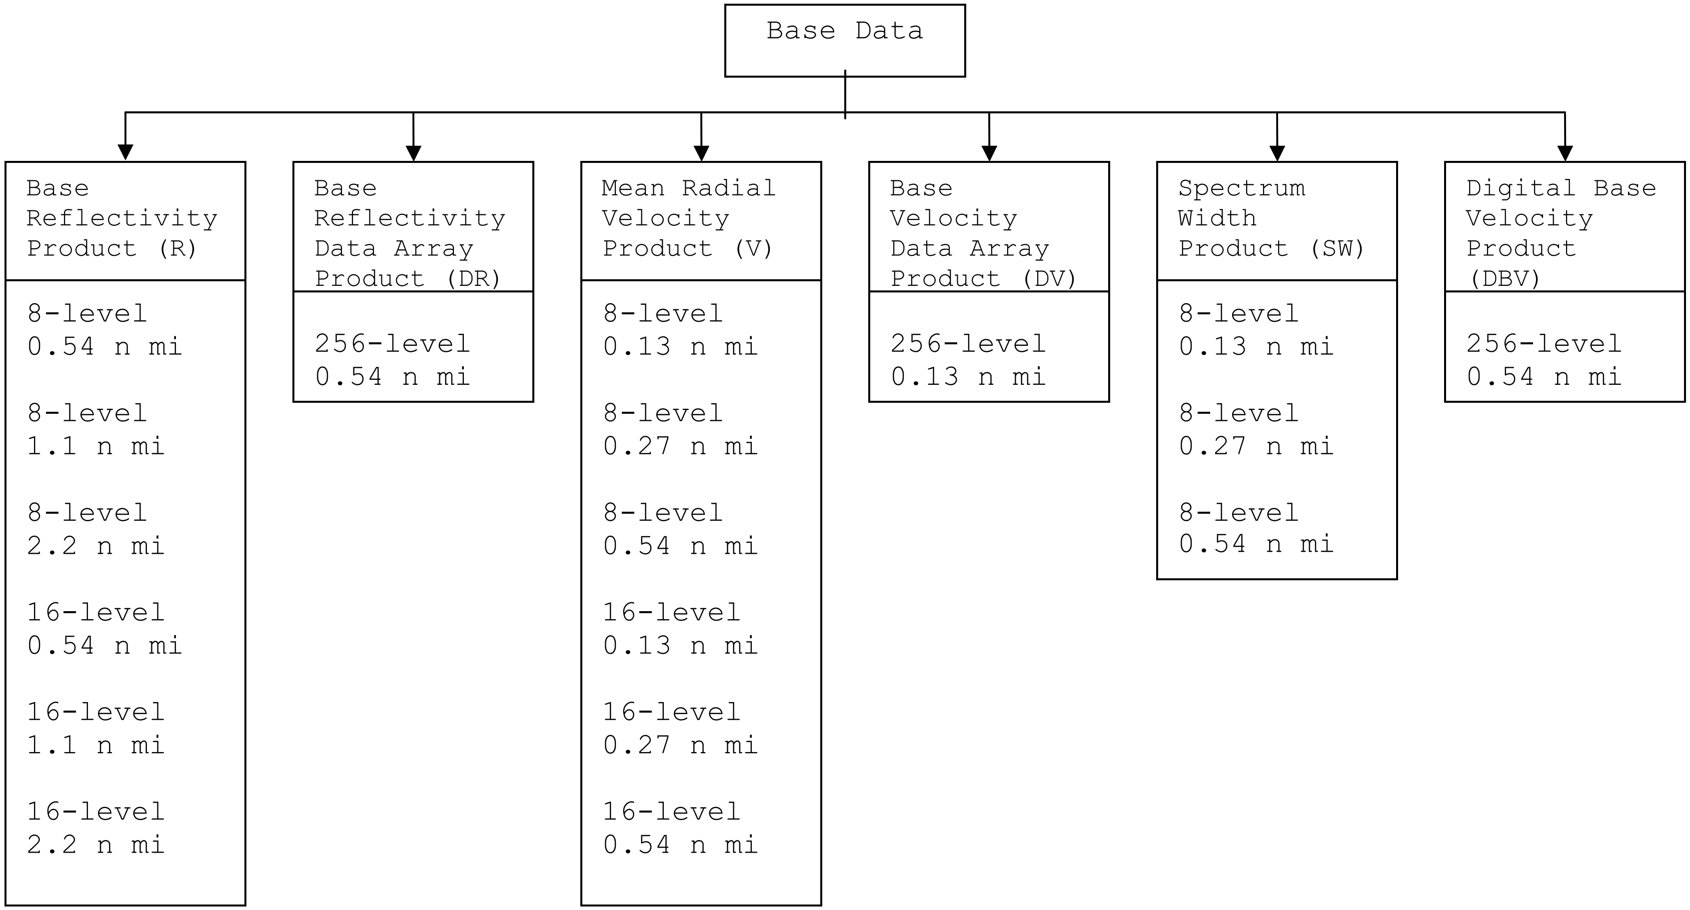

Figure 3-4e
Flow of Data for Generation of Products (Continued)

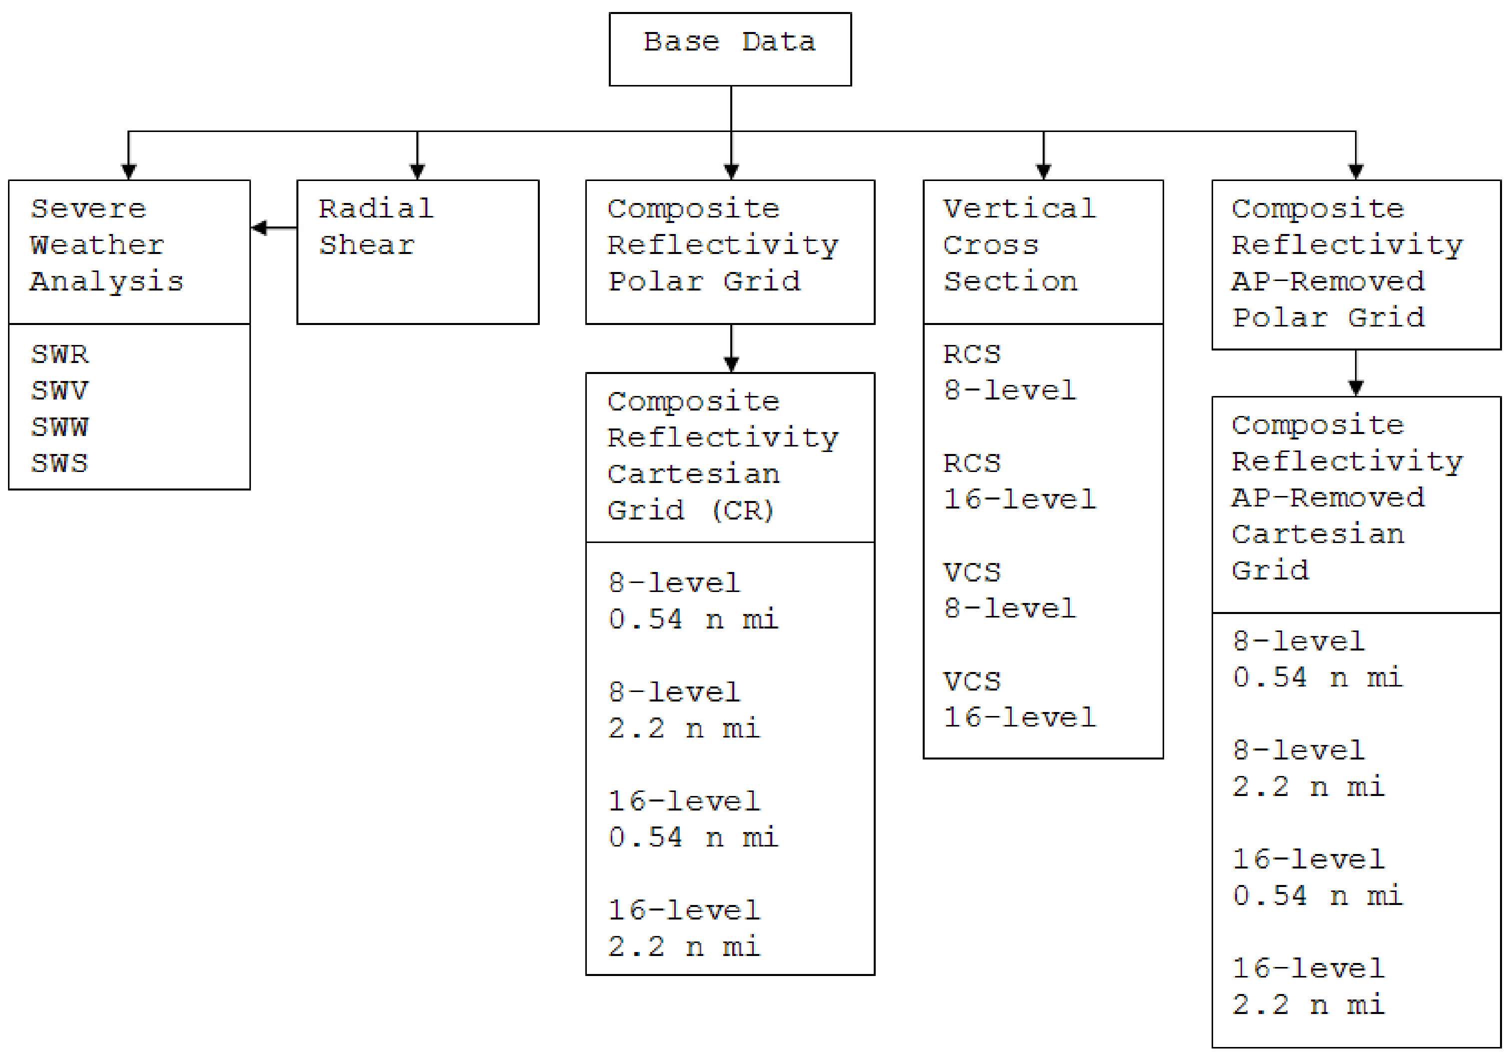

Figure 3-4f
Flow of Data for Generation of Products (Continued)

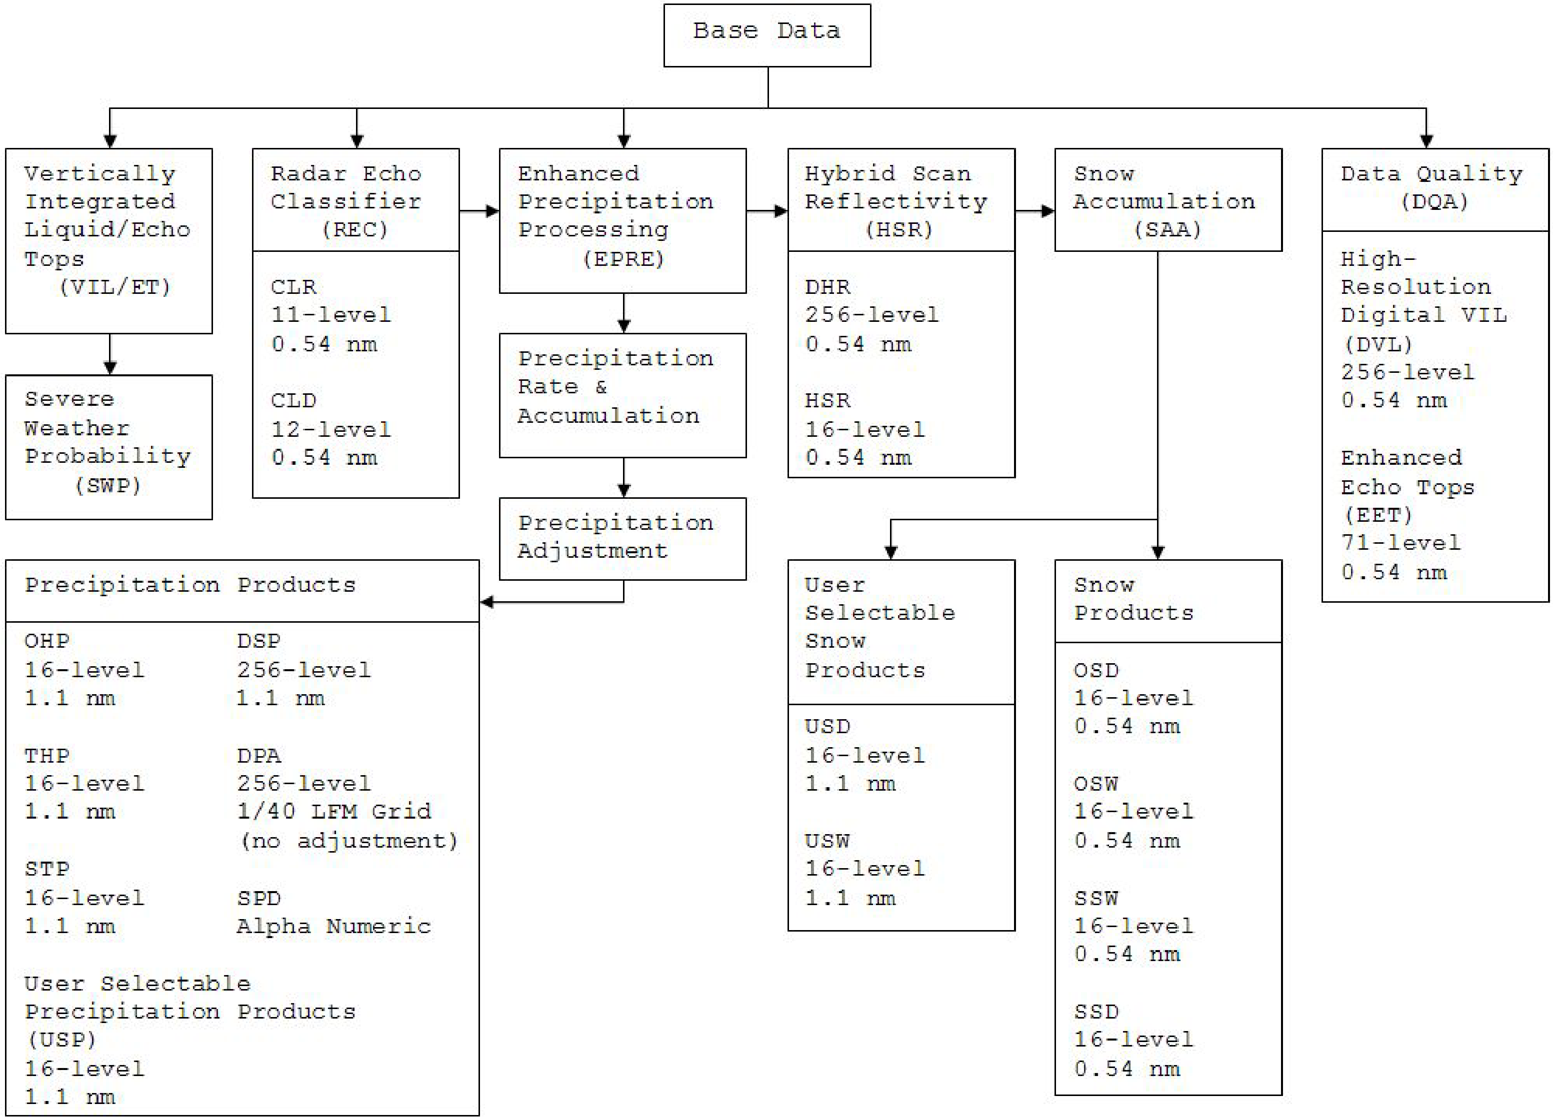

Figure 3-4g
Flow of Data for Generation of Products (Concluded)

Mean Radial Velocity and Spectrum Width. These tasks obtain buffers for the required number of products to be generated (up to six each). The task processes only the number of bins required for the range (0.25 km (0.13 nm) x 1° resolution needs ~ 59 km (32 nm) of bins). For the 0.25 km (0.13 nm) x 1° versions, every bin of data is included in the product. For the 0.5 km (0.27) nm x 1° and the 1 km (0.54 nm) x 1° versions, only every second and fourth bin, respectively, are used for product generation. The only other data processing is run length encoding the bins of data. These products can be produced for any or all elevation cuts within a VCP.

Composite Reflectivity. The Composite Reflectivity polar grid task determines the polar grid coordinates for each radial and for each bin within the radial. The grid value is replaced if a higher value is determined at a higher elevation scan corresponding to the grid value in question. After the grid is completed, the Composite Reflectivity data are mapped from polar grid coordinates to Cartesian grid coordinates. When this process is completed, the Composite Reflectivity Cartesian grid task determines which holes obtained during grid conversion need to be filled. For each hole, a search of neighboring bins is used. If at least two adjacent bins contain data values above threshold, the hole is filled with the average of these values. The final process is run-length encoding of the data for distribution.

Severe Weather Analysis. Radial base data and radial shear data (obtained from the High Resolution Shear Algorithm Output) are input to the task. The only processing is run-length encoding of the data within the defined window. These products can be produced for any or all elevation cuts within a VCP.

Vertical Cross Section. This task builds a Cartesian map of radial data bins required for the points selected. At the end of the volume scan, holes that have data above and below, and are not range folded or below threshold, are filled by interpolating downward. The data are then run-length encoded for distribution.

8-bit Reflectivity. This task builds a reflectivity product formatted as a data array with the same spatial resolution and same data quantization as the input base data. This product can be produced for any or all elevation cuts within a VCP.

8-bit Velocity. This task builds a Mean Radial Velocity product formatted as a data array with the same spatial resolution and same data quantization as the input base data. This product can be produced for any or all elevation cut within a VCP.

3.3.4.3.2 Products Produced from Hydrometeorological Algorithms. The remainder of the products are generated from either radial data and algorithm outputs or from the Output of other algorithms. Details of each algorithm's processing are contained in Part C of this Handbook. All algorithms and tasks implemented in the RPG have an Output. In some cases, these outputs may be in the form of a product. In other cases, the Output is used as an input to another algorithm or task, such as the storm segments output, which is the sole input to the Storm

Components Algorithm. Other algorithms and tasks, such as hail and storm structure, produce outputs for products along with storm attribute information for use by the Mesocyclone Detection Algorithm, Storm Relative Mean Radial Velocity task, and the Combined Attributes Table, which, in turn, passes the storm attribute information to alert processing and the Radar Coded Message. Some algorithms have outputs that are written to an intermediate file and are not currently used in production of a product.

3.3.4.3.3 Products Produced from the Outputs of Other Products. As noted in Section 3.3.4, there are products that are composed of the outputs from other products. These products and those that contribute to their composition are:

Product	Contributing Products
Combined Attributes Table	Hail, Storm Structure, M, MD, TVS, STI
Radar Coded Message	Combined Attributes Table CR Polar Grid Velocity Azimuth Display
User Alert Message	Combined Attributes Table CR Cartesian Grid One-Hour Precipitation Accumulation Severe Weather Probability Velocity Velocity Azimuth Display VIL/Echo Tops.

3.3.5 Product Distribution. After being generated, products are stored temporarily, on disk, in a product database. The distribution function reads the products from the product database to satisfy user requests. All products are required to be stored for at least two volume scans in the event product transmission delays prevent some products from being distributed immediately after generation. Furthermore, the database is sized large enough to accommodate storing products for up to 6 hours to satisfy user one-time product requests.

3.3.5.1 Temporary Storage and Product Storage Loadshedding. After a product is generated, the product is stored within a product database. The amount of time the product resides in the database depends on the storage time specified in the RPG Product Generation Table and whether or not Product Storage Loadshedding occurs. If no storage time is specified, the product will be stored for a minimum of two volume scans. If a storage time is specified for a product type, the minimum time can be as short as 30 minutes or as long as 360 minutes. See Figure 3-3 for the types of RPG loadshedding.

The product database is sized to hold at least 6 hours of products in a full load scenario. When a product is generated, the product is marked with an expiration time which is the sum of the product generation time and storage time. When a product has been stored beyond its expiration time, the product is marked for expiration. For Product Storage Loadshedding purposes, the current utilization is the ratio of the number of products in the database not marked for expiration to the total number of products the database can hold. Like with all loadshed categories, there is a loadshed Warning and Alarm Level (specified as a percent) for Product Storage. If the number of products in the database that are not marked for expiration exceeds the Alarm Level, then a sufficient number of products are marked for expiration in order to reduce the current utilization down to the Warning Level. This event triggers the Product Storage Loadshed Alarm.

When a product is marked for expiration, is it still available in the product database. Only when the product database is full (the total number of products in the database, either marked for expiration or not, reaches the maximum), is an older product removed to make room for the newer product. The product that is removed is one that has been marked for expiration.

3.3.5.2 Product Distribution and Product Distribution Load Shedding. After a product is generated, the Product Distribution module is notified that the product is ready for distribution. This module coordinates the distribution of all data (products, messages, and alerts) to External User and display systems.

After being notified that a product has been generated, the Product Distribution module examines each External User request list. For each list containing the product, the Product Distribution module places information about the product in the distribution queue corresponding to the External User. The position in the queue is based on transmission priority and the product priority defined in the Product Attribute Table. The transmission priorities (highest to lowest) are defined as follows:

5. Alert Messages
4. Product Request Response and General Status Messages
3. Alert products and User Alert Messages
2. One-time products and Routine products (High Priority)
1. Routine products (Low Priority)

Product priority is a number between 1 and 100 and varies from product to product. The value assigned to a particular product depends on the perceived value of the product to the user community, with higher numbers denoting higher priority or value products.

A product queued for transmission on a communications line is transmitted based on its position in the distribution queue. That is, higher priority products are transmitted prior to lower priority products. Consequently, lower priority products could conceivably wait in the distribution queue long periods of time on heavily utilized communications lines.

Distribution load shedding is based on length of time in the distribution queue for a particular communication line and prevents long transmission queue wait times. When a new product enters the distribution queue for a communication line, a check is made to see whether any product has been sitting in the queue for a time exceeding the product "Alarm_level * volume_time."

If true, then products are shed until there are no products that have been waiting in the distribution queue for a time exceeding the product "Warning_level * volume_time."

Distribution load shedding attempts to remove products from distribution depending on age or length of time in the distribution queue with older products (and, consequently, those of lower priority) shed first.