User Guide

IRIS Radar

IRIS Product and Display

VAISALA

Table of Contents

List of Figures

List of Tables

1. About This Document

1.1 Version Information

This document explains how to use IRIS Radar products.

Table 1   Document Versions

Document Code	Date	Description
M211319EN-G	November 2017	This manual. Seventh version.
M211319EN-F	May 2017	Sixth version.
M211319EN-E	April 2017	Fifth version.

1.2 Related Documents

Table 2   Weather Radar Documentation

Document Code	Name
M211315EN	*IRIS and RDA Software Installation Guide*
M211318EN	*IRIS Programming Guide*
M211316EN	*IRIS and RDA Utilities Guide*
M211319EN	*IRIS Product and Display Guide*
M211317EN	*IRIS Radar User Guide*
M211452EN	*IRIS and RDA Dual Polarization User Guide*
M211322EN	*RVP900 Digital Receiver and Signal Processor User Guide*
M211320EN	*Radar Control Processor RCP8 User Guide*

For information on changes made since your current release was installed, download the latest document versions and check the IRIS and RDA Release Notes from www.vaisala.com.

Vaisala encourages you to send your comments or corrections to helpdesk@vaisala.com

1.3 Documentation Conventions

WARNING! **Warning** alerts you to a serious hazard. If you do not read and follow instructions carefully at this point, there is a risk of injury or even death.

CAUTION! **Caution** warns you of a potential hazard. If you do not read and follow instructions carefully at this point, the product could be damaged or important data could be lost.

Note highlights important information on using the product.

Tip gives information for using the product more efficiently.

Lists tools needed to perform the task.

Indicates that you need to take some notes during the task.

2. IRIS Product Overview

IRIS is an advanced hardware and software product for Doppler and non-Doppler weather radar applications.

IRIS provides features required for the operation of a radar network and distribution of radar products, including local and remote radar control, signal processing, product generation, and display.

For more information, see *IRIS Radar User Guide*.

The following tables provide an overview of the products supported in IRIS. The optional products require additional licensing.

Table 3   IRIS Products

Product	Description
BASE Echo Base	**BASE** is used to determine the base of echoes.
BEAM Antenna Beam Pattern	**BEAM** is a full screen cross-section format image showing range-averaged intensity in azimuth and elevation coordinates. **BEAM** is used during calibration and alignment and to verify antenna patterns.
CAPPI Constant Altitude PPI	**CAPPI** (Constant Altitude PPI) is a horizontal cut at a selected altitude used for surveillance and severe storm identification. It is also useful for monitoring the weather at specific flight levels for air traffic applications.
FCAST Forecast	**FCAST** is an array of direction and intensity vectors used to shift a current weather display interactively.
GAGE	The **GAGE** product stores reports from rain gage sensors for display and use in product generation. These products are produced using an input pipe from data sent by a weather station system. Because **GAGE** products are not produced from radar data, they do not need to be configured or scheduled.
HMAX	**HMAX** displays the height of the maximum reflectivity above each output pixel. This product requires a volume scan.
LAYER	**LAYER** can compute layer averages of any polar data types in the ingest files. **LAYER** can also convert to liquid first and compute `VIL Density`. When computing `VIL Density`, the output is in g/m**3.
Lightening	IRIS can import lightning data from a dedicated Vaisala lightning detection network or from the GLD360 Global Lightning Detection system.
MAX Maximum Reflectivity	**MAX** shows the maximum reflectivity over each pixel as well as the East-West and North-South maximum projects in side panels.

Product	Description
MLHGT Melting Level Height	**MLHGT** displays a map of the melting layer altitudes.
PPI Plan Position Indicator	**PPI** is a full screen image used primarily for weather surveillance purposes.
RAIN1 Hourly Rain Accumulation	**RAIN1** is hourly rainfall accumulation. You can select the Z–R relationship.
RAIN-N N-Hour Rain Accumulation	**RAIN-N** is rainfall accumulation of the last N hours, where N is selected by the user.
RAW Raw Data Product	A **RAW** data file that contains the raw signal processor output parameters (Z, ZT, V, W, ZDR, and so on) in polar coordinates. A *RAW* file contains all the information collected during a measurement task. It is compressed so it can be sent over the network to IRIS Analysis. While **RAW** data cannot be displayed, it is useful for recording on tape/disk/DVD or transferring to another IRIS network computer for product generation.
RTI Range Time Indicator	**RTI** displays time along the horizontal axis and the vertical axis displays range from the radar. Often used for manual scans when observing a fixed target.
RHI Range Height Indicator	**RHI** is a full screen image showing the detailed cross-sectional structure of a storm, used for identifying severe storms, hail and bright band.
SRI Surface Rainfall Intensity	**SRI** (Surface Rainfall Intensity) provides input for the **RAIN1** product to obtain the best possible estimates of accumulated precipitation even at longer ranges from the radar.
STAT Radar System Status	A **STAT** file contains a report indicating the state of all key components of the radar software.
THICK Echo Thickness	**THICK** shows the thickness of cloud echoes. **THICK** is the same as the difference between the **TOPS** and **BASE** values. The **THICK** product also computes the average reflectivity in the layer identified by the selected **dBZ Contour**.
TOPS Echo Tops Map	**TOPS** is a color-coded contour map of the top of a selected dBZ level. Either Z or ZT can be used as the basis for the estimate.
TRACK Interactive Storm Tracking and Forecasting	**TRACK** products show the projected motion for storm features (centroids) based on a series of input products from different times. **TRACK** is an overlay product created in the **Quick Look Window**.
VAD Velocity Azimuth Display	**VAD** is a display of the mean Doppler velocity at a given range as a function of the azimuth angle as the radar antenna rotates through an azimuth scan at a constant elevation.
VIL Vertically Integrated Liquid	**VIL** is a color-coded map of the estimated depth of water (in mm) contained in a selectable atmospheric layer. This is an excellent indicator of severe storms. Either Z or ZT can be used as the basis for the estimate.

Product	Description
VVP Velocity Volume Processing	**VVP** (Velocity Volume Processing) provides line graphs or time against height cross-sections of wind speed, wind direction and divergence against height. You can select the range of influence and the maximum height.
WARN Warning/Centroid	**WARN** is automatic warning and centroid plotting. Automatic warnings can be set for protected areas and user-selectable warning criteria. Output is a warning message and a situation overlay showing the centroid locations of storm features, such as high **VIL** or reflectivity.
WIND Wind Speed and Direction	**WIND** displays wind speed and direction with either wind bards or wind strings. You can specify the range and height of the data, and the range and azimuth spacing of the lines that are displayed.
XSECT Cross Section	**XSECT** is a volume scan product that shows the height cross-section along a user-definable line.

Table 4   Optional IRIS Products

Product	Description
CATCH Cross Section	**CATCH** calculates the precipitation accumulation in subcatchment areas such as watershed areas. This product requires the Hydromet option.
COMP Multi-Radar Site Composite	**COMP** produces a single composite image of radar data from many radar sites.
DWELL Composite in Time	**DWELL** is comprised of a product and composites successive images of a product in time. Moving targets show a "streak" on the display.
NDOP Multiple Doppler	**NDOP** provides a mechanism for including velocity (Doppler) data from another radar site now enabling the measurement of horizontal wind in addition to vertical wind.
SHEAR Wind Shear	**SHEAR** detects wind shear in the atmosphere, allowing the detection of microbursts, gust fronts, mesocyclones, cold fronts, and atmospheric waves.
SLINE Shear Line (frontal boundary)	**SLINE** (shear line, or frontal boundary) marks the transition between two air masses on the image.

3. Configuring IRIS Products

3.1 Configuring Products

IRIS Radar products provide information that can be used directly for weather nowcasting and forecasting. You configure the products in IRIS Radar.

IRIS products can display radar data in many ways. For example:

- The **CAPPI** product shows the distribution of a radar parameter, such as reflectivity or spectrum width, at a constant altitude.
- The **TOPS** product shows a color contour map of the height of a selected reflectivity surface.
- The **RAIN1** and **RAINN** products show the accumulation of precipitation over selectable time periods.
- The **WARN** product checks other IRIS products to determine if there is significant weather, such as wind shear greater than a threshold amount.

Each product is associated with a configured task. Tasks collect information from signal processors and store the data on disk in ingest files.

For more information, see *IRIS Radar User Guide*.

More Information

‣ QLW Live Action Tool: Product Generation and Display (page 215)

3.2 Product Configuration Menu

Most **Product Configuration** menus have the same general format:

Task Summary
> Contains information about the task associated with the product.

Map Projection
> Specify geographical map projection, if desired.

Product Parameters
> Specifies the data, range, and other product-specific options.

Display Parameters
> Select your default color scale, which can be overridden at display time.
> In case of `HClass` data types, select the classification method projected in the
> product. This cannot be overridden at display time.

▶   1. Select **Menus > Product Configuration**.

3.2.1 Associating Products with Tasks


```
                      TASK SUMMARY
TASK Name    │PPIVOL_A    │  Sub TASK │AB│        Max Range │160.5  │
Scan Mode    │PPI         │  DSP Data │2Z T V W SQI         │
Angle List   │Az:Full Circle     El:7 angles from 0.5 to 7.0        │
```

Each product is associated with a task, which provides the ingest data for the product.

TASK Summary information, taken from the associated task, shows the key task configuration parameters.

When associating a task with a product, the system must consider *matched resolution sampling* and constructing *volume scans* for a product.

▶ 1. Select **TASK Name** to show a list of tasks.
Note that certain product types require different types of task associations.

Product Type	Description
Running a product from ingest data	Enter a task name, including wildcard characters. The question mark (?) matches a single character; the asterisk (*) matches any string of characters.
Volume scan product from hybrid tasks	In the task name section: • Specify the input tasks with an underscore and single letter. • Specify the sub-task suffix letters in the **Sub Task** box. Type a ″*″ to indicate all subtasks. • Type individual letters such as **ABC**, **AB**, **BC**, or **AC** to select which portions of the hybrid to include. • Type ″–″ to indicate a range of letters, for example ″A–DF″ means **ABCD** and **F**.
RAW products	Always made from a single task. You must make a separate **RAW** product for each sub task. If you are retrieving **RAW** data from tape or receiving a **RAW** product over the network, there is no task configuration file on your disk. In this case, type the task name directly in the field, exactly as it appears in the **Ingest Summary** menu or the **Tape** menu.
RHI products	Include **RHI** in **RHI** task names to distinguish them from **PPI** scans.

2. To optimize product appearance and system performance, define the matched
    resolution sample.
    For best results, the range bin spacing and number of bins in the task should match
    what is needed for the product. This is called a *matched resolution sample*. The match
    does not need to be exact because IRIS uses interpolation algorithms.
    For best performance, match the range bin spacing, number of bins and maximum
    range in the task configuration, to the desired pixel resolution, pixel scale and
    maximum product range in the product configuration.
    For example, if you are making a low resolution product (240 × 240) to a range of
    120 km (74.6 mi), and you want a single pixel to represent 1 km (0.6 mi) of data,
    configure the task configured with at least 120 range bins spaced 1000 m (3280 ft
    10 in) apart.
    Note that 170 bins can be displayed to the corners of the product. For a medium
    resolution product with 1 km (0.6 mi) bin spacing, the best match is for tasks with 240
    to 340 bins spaced 1 km (0.6 mi) apart. Note that IRIS can make products for any range
    bin spacing that can be specified in the **TASK Configuration** menu, for example, 500 m
    (1640 ft 5 in) or 250 m (820 ft 3 in) spacings can be used as well.

3. If needed, construct a volume scan.
   Some products require volume scan tasks, either **PPI Full** or **PPI Sector** scans taken at multiple elevation angles. Volume scan products include:
   - **BEAM**
   - **CAPPI**
   - **MAX**
   - **TOPS**
   - **VIL**
   - **XSECT**
   - **VVP**
   - **WIND**

For these products, multiple elevation angles are required in the associated task. There is a tradeoff between the number of elevation angles, the quality of the product and the task scan time. More elevation angles produce higher quality products at the expense of taking more time to complete the volume scan.

A typical volume scan is shown in the following figure. In this example, the height resolution is 1 km (0.6 mi) at 60 km (37.3 mi) range, for heights less than 10 km (6.2 mi). A 1° beam is 1 km (0.6 mi) wide at 60 km (37.3 mi), so this scheme matches the antenna resolution. If close range work is important, you must add higher elevation angles to cover the upper regions.

Figure 1   Example of 15-tilt Volume Scan

When constructing a volume scan task configuration, it is a good idea to do a drawing such as the one in the example, with a 1:1 vertical:horizontal scale so that you get a true picture of the sampling geometry.

> The example is corrected for earth curvature, as are all IRIS products.
> Also, because of beam widening effects, the accuracy of all products degrades with range. For example, the beam width at 120 km (74.6 mi) range is 2 km (1.2 mi) across for an antenna with a 1° beam. This is a fundamental limitation of radar sampling.

3.2.2 Configuring Map Projections

You can specify the type of map projection that is used for display. This is useful for combining data from other sources or for generating displays that are not centered on the radar.

Composites, for example, must use a projection so that the data from several radars can be mapped to a single display.

Table 5   Supported Map Projections

Projection	Description[1]
Azimuthal Equidistant (AED)	Choose this for a standard radar display map where azimuth lines are straight and angles are not distorted. This projection has the property that the distance in the map horizontal direction is the same as the distance in the map vertical direction. This means that lines of constant azimuth (radar rays) are straight and circles can be used to represent lines of constant range. Because of this, the AED projection is particularly convenient for radar applications. For radar applications, the radar is usually placed in the center of the map and the radar location serves as the map reference point. Note that lines of latitude and longitude are not straight in this projection. (Page 191)
Equidistant Cylinder	This is used for the digital terrain maps in IRIS (for example, inserted in the **Overlay** utility). (Page 90)
Gauss Conformal	Same as Universal Transverse Mercator except omitting the 0.9996 scale factor. (Page 48)
Gnomonic	All great–circle arcs are straight lines. (Page 164)
Lambert Conic	Lambert Conformal Conic, scale is true along 2 standard parallels. (Page 104)
Mercator	Choose this for true Mercator projection mapping where latitude and longitude lines are straight and intersect at right angles. Additionally, a line of constant bearing from a point is straight so it is useful for navigation. Mercator has the advantage that it is a known standard, so it is easier to combine data from different sources onto a Mercator projection. It has the disadvantage that at high latitudes, the horizontal and vertical scales are different (large distortion). (Page 38)
Perspective	This is often used for satellite images. We assume that the satellite is at geosynchronous height. (Page 169)
Polar Stereographic	Useful near the North Pole. This is essentially the AED projection with the reference point at the North Pole. (Page 154)

Projection	Description[1]
Universal Transverse Mercator (UTM)	This is a projection just like Mercator with the axis reversed at a selectable longitude. The horizontal dimension is compressed by a factor of 0.9996. By convention the reference longitudes should be multiples of 6°. This is used by the military, and is good for regions with long North-South extent. (Page 48)

1) *Page numbers refer to the projection equations from Map Projections – A Working Manual, U.S. Geological Survey Professional Paper 1395.*

▶ 1. To configure a standard radar display centered on the radar:
- For the projection type, select **Projection Configuration > Azimuthal Equidistant**.
- For the **Projection Name**, select **<NONE>**.

> ℹ️ Most of the time you do not need to create or use special projection files.
> If you want the radar off-center, you must configure a projection file as described starting from step 3.
> If you have already configured and saved a projection, go to step 2.

2. To use a pre-configured projection, select **Projection Configuration > Projection Name**.
- A list of available projection files on the system is shown.
- In the example, the radar **Max Range** and **Output Pixel Resolution** are desensitized since they are defined by the projection file that is selected.
- Select **<NONE>** for the **Projection Name** in which case the radar is defaulted to be at the center of the screen. Additionally in the case of an AED projection, the radar is defined to be the map reference point. In this case, the **Max Range** and **Resolution** fields are configurable in the **Product Configuration Menu**.

3. To configure your own projection, select **Setup > Projection Configuration**.
   The **Projection Configuration Menu** opens.
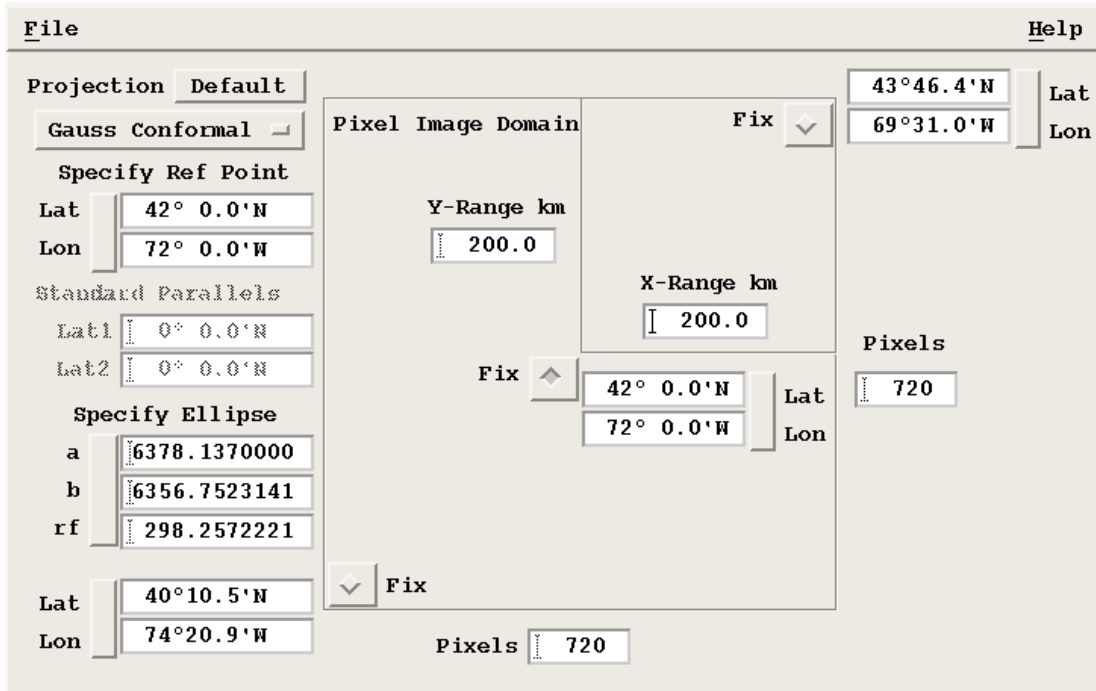
4. Select the **Projection** to define the projection type.

> ℹ️ You must first define the projection. The list of available projection files under **File** is given only for the selected projection type.

5. Select **File > Open** command to give a list of existing Mercator or AED files depending on the selected projection.

6. Select **Reference Point**.
   • For AED projections: The reference point is the location from which all azimuths display as straight lines. Usually set to the radar location.
   • For Mercator projections: The map cannot cross the line 180° away from this longitude, so set it near your location. The latitude defines the range scale.
   • For Polar Stereographic: The is the longitude line which you wish to have vertical in the map. Effectively this rotates the map. Again the latitude defines the range scale.
   • For UTM: Only a longitude is selected. This is the center meridian of the UTM strip. By convention these are spaced every 6°, starting at 3°. Use a value near your longitude.
   • For Lambert Conic: Also define 2 standard parallels.

7. Select **Save** to name your projection.
   The name of the open projection file is displayed in the menu title bar (*AED_200_KM* in the example).

8. Fix a corner or center location.
   To configure a projection, you must know the latitude and longitude of either the NE or SW corner or the center of the display.
   Select **Fix** at the location that you know and then input the LAT/LON coordinates.
   With **Fix** selected, you can change other parameters of the projection geometry without losing you coordinate information.

9. Specify the projection coverage area (region of the projection) in one of 2 ways:
   - Specify the latitude and longitude of one of the 2 unfixed points.
   - Specify the distance North (**Y-Range**) and the distance East (**X-Range**) from the center.

   The projection menu acts as a spread sheet. When you change a parameter, the other parameters adjust to reflect the change.

   > The spreadsheet algorithm converges on a correct solution even for strange projections. One example where it may not converge is for an AED projection where the **Reference Point** is several thousand kilometers from the center of the projection region.
   > If you encounter a problem with the spreadsheet, check the reference point and select **Default** to load some valid numbers.

10. Specify the pixels.
    The number of pixels that are used to represent data in the projection is specified for the X- and Y-directions of the output display.
    If you have specified the other aspects of the projection, input the correct number of X-pixels (horizontal direction). The number of Y-pixels adjust automatically.
    If you adjust the number of Y-pixels, then other parameters of the spread sheet change, such as the LAT/LON of the unfixed points and the **Y-Range**.

11. Fine-tune the spread sheet.
    Since all of the projection geometry parameters are linked through the spread sheet algorithm, when you change one parameter, other parameters adjust automatically. This means that you may need to compromise. The spread sheet makes it easy to experiment with different compromises.
    For example, if you want a 480 × 480 pixel Mercator display, both the X- and Y-ranges cannot in general be made equal (for example, 100 km (62.1 mi)). In this case, the compromise can be, for example, an X-range of 100 km (62.1 mi) and a Y-range of 99.5 km (61.8 mi) to get a perfect 480 × 480 display.

3.2.3 Product Parameters

In **Product Parameters**, define the data taken from the task and how to display the data.

```
          PRODUCT PARAMETERS
Data:Display    [ dBZ
  Max Range     ][160.0
   EL Angle     ] 0.5
 Clip Range     ]  0.0 km
Clip Height     ]  0.0 km
ZR relation     ][200**1.60
XY Smoother     ] 0.2
```

Data : Display

Select **Data:Display** to show a list of valid display parameter options. The available choices vary depending on the product type.

This parameter specifies from which data parameter the product is derived. For example, Z, V, W, ZT, or ZDR. This information comes from the task associated with the product. Some data parameters can be displayed in more than one way.

Table 6   Data Types

Data type	Description
Ah, Av	Integral attenuation for horizontal (H) and vertical (V) channels.
Azdr	Integral attenuation of ZDR (dB) format.
CSR	Doppler channel clutter-to-signal (CSR) ratio of dBT to -dBZ.
dBT, dBTh, dBTv, DBTr	Total power.
dBZ	Clutter corrected reflectivity
dBZt	Uncorrected reflectivity
HCLASS	Hydrometeor classification Estimated hydrometeor type in the precipitation area.
KDP	Specific differential phase An indicator of the rate of change of the phase difference between horizontally and vertically polarized pulses of the radar. A greater horizontal shift results in a positive KDP value, and a greater vertical shift results in a negative KDP value. Typical cause for a high KDP area is heavy rain.
LDRH, LDRV	Linear Depolarization Ratio H to V (or V to H). The ratio of cross-polar to co-polar reflectivity measured in dB.
LOG	Log receiver signal-to-noise ratio.

Data type	Description
PHIH, PHIV	Horizontal (V) or vertical (V) differential phase
	Phase difference for the total round trip between radar and the volume where the signal is reflected. PHIH is measured between HH and HV channels. PHIV is measured between VV and VH channels.
PHIDP	Differential phase
	The phase difference due to propagation between the HH and VV channels of the radar.
PMI	Polarimetric meteo index.
R	Rate of accumulation of precipitation in units of mm/hour. For snow, this is usually refers to the liquid equivalent.
RHOHV, RHOH, RHOV	Correlation coefficient between HH and VV (or HH & HV / VV & VH) channels)
	Higher (>0.95) values indicate uniform precipitation areas and lower values more mixed hydrometeor types, such as melting snow, wet snowflakes, or airborne debris.
SNR	Signal to Noise Ratio
	Generic measurement of signal-noise ratio in dB.
SQI	Signal Quality Index
	A value between 0-1 that measures the signal's Doppler coherency, that is the correlation between the signal and its Doppler lag.
	• 0 indicates white noise • 1 is the perfect Doppler point target
T	Total Reflectivity
	Total power returned to the radar in reflectivity units. It typically represents the horizontal reflectivity without ground clutter correction.
TV, TE	Total Vertical (HV Enhanced) Reflectivity
	Total reflectivity from the vertical polarization channel (TV) and combination of the horizontal and vertical channel (TE).
V	Velocity
	Average radial velocity (towards or away from the radar) of detected hydrometeor areas.
VC	Corrected Velocity
	Same as Velocity, but corrected for effects of range folding and velocity folding.
V: SHEAR, Vc: SHEAR	Velocity and corrected velocity of wind shear.
W	Spectral Width
	Variability of Doppler velocity values within the measurement area.
XCOR	Polar cross-correlation, uncorrected rhohv. Because this value is not noise corrected, it is a direct indicator of the PHIDP uncertainty
Z	Reflectivity
	Usually referred to as dBZ, this is the common data type that measures radar signal reflectivity, and is used to estimate precipitation intensity from that.
	All Z measurements are corrected for ground clutter.

Data type	Description
ZV, ZE	Vertical (HV Enhanced) Reflectivity Total reflectivity from the vertical polarization channel (`ZV`) and combination of the horizontal and vertical channel (`ZE`). Corrected for ground clutter.
ZC	Corrected Reflectivity Same as `Z`, but corrected for attenuation and beam blockage effects.
ZDR	Differential Reflectivity The ratio of SNR in the horizontal channel to the SNR in the vertical channel. Positive values indicate more prominent horizontal echoes and negative values more prominent vertical echoes. Larger hydrometeor sizes are usually identified by high positive ZDR values.
ZDRC	Corrected Differential Reflectivity Same as `ZDR`, but corrected for attenuation and beam blockage effects.

Product Maximum Range

> If you select a custom **Map Projection Name** (other than `<NONE>`), you cannot configure this field in the product output menu. This is because the maximum range is determined by the projection configuration.
> See 3.2.2 Configuring Map Projections (page 21).

This is the maximum range configured for the product, which must be less than or equal to the data range. The maximum range is the range in the E-W / N-S direction centered on the radar. This means that the range to the corners of the resulting display is greater.

ZR Relation

Weather radars do not measure the rainfall rate (`R`) directly. Instead, they measure the equivalent radar reflectivity factor (`Z`).

In general, the larger the reflectivity factor, the larger the rainfall rate. The `ZR` relationship is a standard technique for estimating the rainfall rate based on the measurement of the radar reflectivity factor. The have the following general coefficient exponent:

```
Z = 200 R 1.6 (the classic Marshall Palmer relationship)
```

Table 7 **ZR Relation** Options

Options	Description
Fixed	The default value configured in **Setup** is first inserted. You can type or change the coefficient and exponent.

Options	Description
Setup	The coefficient and exponent values set in **Setup > Product > Product Generation** are retrieved when the product runs. If the setup values change, the new values are used.
Disdrometer	The coefficient and exponent values are derived from a **GAGE** product with the assigned name **DISDROMETER**. The **GAGE** product contains (optionally) the coefficient and exponent derived from a disdrometer that measures the size distributions of particles. When the product runs, IRIS looks for the closest **DISDROMETER** product, first searching in time (± 1 hour) and then opening the file and searching in space by comparing LAT/LON (within 200 km (124.3 mi) is valid). If a valid coefficient and exponent are found, they are used in the product generation, otherwise the **Setup** defaults are used. Only available with the IRIS/Hydromet option.

XY Smoother

The **XY Smoother** can significantly improve the appearance of products for presentation and improve data compression. The smoother performs image processing to average the color contours and interpolate over small gaps or missing pixels in the final Cartesian image.

Specify the length of the smoother in km.

CAUTION! Smoothing requires substantial CPU resources. The internal limit on the smoother length is 60 pixels, much longer than you probably want to use.

To produce many smoothed products in a regular operation, Vaisala suggests the following:

- Use low resolution products together with smoothing.
  It is faster to make a low resolution product and smooth it than to make a medium resolution product and smooth it.
  In many cases, the results are nearly identical, so there is no benefit in smoothing high or medium resolution products.
- Use the minimum smoothing length to get the desired effect. Do not over-smooth and do not use a length that is larger than you need.
  Larger smoothing lengths require more computation.

3.2.4 Display Parameters

The display parameters determine how the product appears.
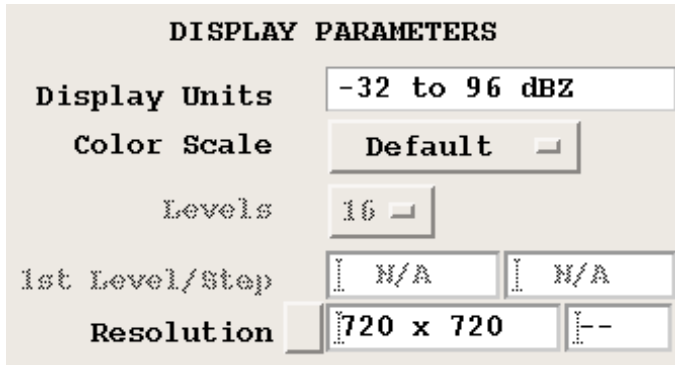
Display Units

Most data types are numerical. **Display Units** shows the range of values and the units for the product's output values.

For example, the Echo Tops algorithm outputs values in km to the nearest 100 m (328 ft 1 in) spanning the range of 0 to the maximum data height defined with the **setup** utility or 25.5 km (15.8 mi), whichever is less.

The following table shows the values that you can assign to the display units.

Table 8   Display Units

Display Parameter	Maximum Display Range, Display Units, or the Classifier method
dBZ	-32 … 96 dB mm^6/m^3
dBZt	-32 … 96 dB mm^2/m^3
Rain	0 … 255 mm/hr or 0 to 10.2 inches/hr
Liq	0 … 1000 mm
Vel	$-V_u$ to V_u m/s
Width	0 … V_u m/s
ZDR	-8 … +8 dB
TOPS	0 … 25.5 km or height configured in **Setup**
VIL	0 … 65.0 mm
WIND	Configured in **Setup**
SHEAR	+-25 m/s/km
HClass	`Meteo`, `Precip`, `Cell`, `Meteo+Precip`, and `Meteo+Cell` The first **Classifiers** field lets you to choose which classifier outcomes are projected in the product. The selector lists the available classification algorithms and their merged combinations.

Color Scale, Levels, 1st Level/Step

The color scale relates numerical data values to colors.

The **Scale** field lets you choose whether to use a uniform scale with a start value and a constant step, or a custom color scale that has been preconfigured for your system:

- Custom color scales are defined with the **Color Setup** utility. See *IRIS and RDA Utilities Guide*.

  Each color scale has a name that has been defined in **Color Setup**.

  Names can reflect scales that are appropriate for different seasons such as `Summer` or `Winter`.

  If in doubt, select **Default** which is usually configured to be a reasonable scale. The number of scale steps are filled-in automatically since this is also part of the custom color scale.

- When you choose a uniform color scale, you pick the number of level steps (2 … 16). You must also set the following fields:
  - The **1st Level** field sets the numerical value used to label the right side of the first level in the color scale for the product.
  - The **Step** field sets the spacing between the color levels. The numbers in the color legend in the product are spaced by this value.

The color scale also associates the classification data to the class descriptors (class identifiers with class legends). Sets of class legends can be customized as **Classifier** sets in your system (with the **HydroClass Name Editor** in the **Color Setup** utility), to match the available classifier methods.

The following table shows the units you use for the **1st Level/Step**, depending on the display parameter.

Table 9   Display Parameter Formats

Display Parameter	1st Level/Step Format
dBZ	+ - XX Whole dBZ
dBZt	+ - XX Whole dBZ
Rain	+XXX.X mm/hr
Liq	+XXX.X MM
Vel	+- X.X m/s
Width	+ XX.X m/s
ZDR	+- X.X dB
TOPS	+ XX.X km
VIL	+ XX.X mm
WIND	NA
SHEAR	+- XX.X m/s/km

> For the mean velocity, selecting a level step of 0 causes the velocity to be displayed for the entire unambiguous interval between $-V_u$ to $+V_u$.
>
> Selecting an odd number of levels produces a legend with a band centered on 0 velocity. Selecting an even number of velocity levels does not produce a legend with a band centered on 0 velocity, but rather a color break exactly at 0 velocity.

When choosing a uniform color scale, you have more flexibility to choose colors. The associations of the class identifiers with class legends is fixed by the **Color Setup**.

The following figure shows 2 examples of the display configuration parameters, one for ZT and one for velocity, and the resulting color legends that would appear with the product.

In the case of ZT, the example shows that the bottom of the scale "thresholds" such that data less than -20 dBZ are not displayed, while the top end "saturates" since the top color includes all values greater than 50 dBZ. When you select uniform scales, the behavior (saturation vs thresholding) of the end of the scales cannot be changed. This behavior is inherited from the custom color scale named `Default`, which is configured in **Color Setup**.

In the velocity example, by entering 0 for level step, the display automatically spans the full Nyquist interval. Note that both ends are set to saturate for velocity.

The format of the color legends, with regard to how the numerical legend labels relate to the color band boundaries, is configured in the **Color Setup** utility for each parameter. This utility offers considerable flexibility for defining how the color legends are constructed for each data parameter. See *IRIS and RDA Utilities Guide* .

Z or ZT Example

Display Units	−32 to 96 dBZ
Number of Levels	8
First Level	−10
Level Step	10

V Example

For unambiguous velocity +−32 m/s	
Number of Levels	9
First Level	NA
Level Step	0

Figure 2   Color Legend Format Example

Product Picture Resolution

> If you select a custom **Map Projection Name** (other than **<NONE>**) you cannot configure this field in the product output menu. This is because the resolution is determined by the projection configuration.
> See 3.2.2 Configuring Map Projections (page 21).

IRIS can produce product pictures in virtually any resolution, from 16 × 16 to 3100 × 3100 pixels. IRIS products are produced specifically for the display resolution that is requested, optimizing the match between the display pixels and the actual radar data. This means that high resolution products are not merely low resolution products with replicated pixels. Likewise, low resolution products are not produced by degrading a higher resolution product. When you request a high resolution product, you get the best possible image that can be generated from the original data. If a low resolution product is requested, the product is computed efficiently by calculating the product only at the requested resolution. In addition to efficient generation, low resolution products can be transmitted over a communications link and displayed more rapidly because they contain fewer pixels.

The **Resolution** field in the **Product Configuration** menus show the X-Y number of pixels. These can be changed by entering different numbers of pixels, or by popping up a menu of the default low, medium or high picture resolutions. The defaults are shown in the following table for both PPI (square image region with legends to the right) and RHI formats (rectangular image region with legends beneath). The defaults are optimized to the resolution of the standard IRIS display devices.

Table 10   Default Picture Resolutions

Default Resolution RHI (Rectangular) Format[1]	PPI (Square) Format[2]			
	X-Pixels	**Y-pixels**	**X-Pixels**	**Y-Pixels**
Low	240	240	288	136
Medium	480	480	600	290
High	720	720	840	530
XHigh	940	940	1060	750

1)   *Rectangular image region with an optional legend beneath.*
2)   *Square image region with an optional legend to the right.*

Regardless of the product type, IRIS determines the best way to display your product in either **PPI** or **RHI** format. If your image does not exactly match the pixel size of a target display device, IRIS adjusts it to make a best fit. For example, if you output a low resolution image to a high or medium resolution printer, IRIS doubles the pixels so that the display fills the image area. Likewise, IRIS shrinks an image if it has more pixels than can be displayed on an output display device.

If you want the impact of the large screen, but you do not want to burden IRIS with creating high resolution products, specify 360 × 360 display resolution. IRIS doubles these numbers to fit exactly in the large display window.

Likewise, 240 × 240 images are tripled to fit exactly in a large display, or doubled to fit in a medium display.

For more information on data formats, see *IRIS Programming Guide*.

3.3 **BASE**: Echo Base Product

Figure 3   **BASE** Example Configuration

BASE is used to determine the base of echoes.

Similar to the echo **TOPS**, the user specifies a **dBZ Contour** for the base. For each output pixel in the product, the algorithm searches downward through successive elevation angles. A base height is determined by linear interpolation when the specified **dBZ Contour** is crossed from greater to lesser value.

The final output of the product is a color-coded map of echo base heights for the selected dBZ contour. Optional Cartesian smoothing can be applied.

> C, S and X band radars detect precipitation particles such as rain or snow. They are not sensitive to cloud particles which are much smaller (of order a factor of 1000 smaller in diameter). The **BASE** product cannot be used to detect the cloud base height for radars transmitting in these bands.
> Millimeter wavelength radars (K and W bands) can detect cloud particles.

> **i** Because of earth curvature, ground clutter and beam blockage, radars are poor at looking for features close to the surface.
> This means that the **BASE** product cannot detect echo bases to very low levels.
> When precipitation in excess of the specified dBZ contour is reaching the ground, or near to reaching the ground, the **BASE** product displays a special color to indicate that there is a **BASE**, but its height cannot be determined.

▶ 1. Select **Type > BASE**.

2. Configure the **BASE** product:

   Data:Display
   Choose from `dBT Height`, `dBZ Height`, or `dBZc Height`

   Max Range
   Select the maximum range for the product. Recall that earth curvature effects are greater at greater ranges.
   You can display a PPI at your lowest elevation angle and use the cursor tool to see what heights are possible to observe at various ranges.
   For example, for a 0.5° elevation angle at 100 km (62.1 mi), the beam is centered at 2.3 km (1.4 mi). It is not possible to detect echo bases lower than this.

   dBZ Contour
   Select the dBZ threshold for the echo base.

3. Apply the **AZ/EL Smoother**.
   Enter the values in degrees, first for the azimuth and then for the elevation direction.
   Typical values are 1.0, 1.0.

More Information

‣ Display Parameters (page 29)
‣ Product Parameters (page 25)
‣ Associating Products with Tasks (page 19)
‣ TOPS: Echo Tops (page 91)

3.4 **BEAM**: Antenna Beam Pattern Product

Figure 4　**BEAM** Example Configuration

BEAM is a full screen cross-section format image showing range-averaged intensity in azimuth and elevation coordinates.

The **BEAM** product is similar to a cross-section in that it shows a quasi-vertical slice of the atmosphere. It is constructed from a **PPI** volume scan. The product is used primarily for system testing and evaluation purposes, such as:

- Antenna beam pattern testing. Typically a radiator is on a tower a few km from the radar. The radar is then scanned in sector or full **PPI** mode around the target using high resolution (for example, <0.5°) spacing in **AZ** and **EL**. The transmitter is typically turned off for this. The **BEAM** product can then be used to display the antenna response to the radiator.
- Calibrated sphere measurement. Similar to above except the radar is radiating. The sector scan is made about the calibration sphere which is typically suspended from a tethered balloon. The resulting display shows the location and returned power from the sphere. This can be used for full system calibration.

At near horizontal incidence (low elevation angles), the **BEAM** product is similar to a cross-section. Since the product can be configured to average a range interval, the **BEAM** product can give a picture of the weather or other targets in a spherical "slab" defined as a range, azimuth and elevation interval. The cross-section on the other hand, attempts to show a planar "slice" through the weather.

1. Select **Type > BEAM**.

2. In **Data:Display**, choose `Z dBZ`, `V V`, `W W`, or `T dBZt`.

3. In **Min/Max Range**, Select the range interval over which averaging is performed. The **BEAM** product performs averaging of `dBZ` or `dBT` values without linearizing. All distances are in km.

> ℹ️ IRIS has a selectable maximum height above which data are not recorded (for example, 20 km (12.4 mi)). When selecting the **Min/Max Range** range, consider that, at far ranges, this height is exceeded in the upper angles of a volume scan.
> For example, if the height cutoff is set to 20 km (12.4 mi), and a volume scan uses elevation angles up to 20°, the maximum range that includes all of the data is 55 km (34.2 mi).

4. In **Min/Max E L**, define the elevation window.
   For example, if the volume scan has elevation angles between and including 0 and 5°, enter the values 0.0 and 5.0.

5. In **CW AZ Interval**, define the azimuth window.
   The first value corresponds to the left side of the display, the second value to the right side of the display.
   From left to right on the display, the angles are in the clockwise direction.
   For example, clockwise interval of 0.0 to 90.0 represents a 90° sector, while the interval 90.0 to 0.0 represents a 270° sector.

6. Apply the **AZ/EL Smoother**.
   Enter the values in degrees, first for the azimuth and then for the elevation direction.
   Typical values are 1.0, 1.0.

More Information

‣ Display Parameters (page 29)
‣ Product Parameters (page 25)
‣ Associating Products with Tasks (page 19)

3.5 **CAPPI**: Constant Altitude Plan Position Indicator

wind   CAPPI Product Configuration: Z_010_120

File   Menus   Type   Commands   **Help**

TASK SUMMARY

TASK Name		PPI_VOL	**DSP Data**	ZV
Scan Mode		PPI Full	**Max Range**	240.0
Angle List		AZ: –NA– EL: 15 Tilts from 0.3 to 34.9		

Map Projection   Azimuthal Eqdist ⌐   **Projection Name**

PRODUCT PARAMETERS   **DISPLAY PARAMETERS**

Data:Display	Z : dBZ	**Display Units**	-32 to 96 dBZ
Max Range	120.0	**Color Scale**	Default ⌐
CAPPI Height	1.0 to 10.0	**Levels**	16 ⌐
CAPPI Fill	☑	**1st Level/Step**	N/A	N/A
ZR Relation	200 ** 1.6	**Resolution**	480 x 480	10
XY Smoother	2.00	**Storage Format**	☑ Data	⬚ Pict

Figure 5   **CAPPI** Example Configuration

CAPPI (Constant Altitude PPI) is a horizontal cut at a selected altitude used for surveillance and severe storm identification. It is also useful for monitoring the weather at specific flight levels for air traffic applications.

CAPPI requires a **PPI** volume scan at multiple elevation angles. The number of angles and their spacing depends on the range and height of the **CAPPI** you want to produce.

CAPPI also supports calculation of **SHEAR** data. The configuration options are the same as the **SHEAR** product, with choices of radial, azimuthal, and elevation shear.

▶ 1. Select **Type > CAPPI** and
      The **CAPPI Product Configuration** menu opens.

   2. In **Data:Display**, specify the data type for you wish to make a **CAPPI**.
      Most IRIS data types are available.
      If you wish to generate rain rate, select which input data type is used for the conversion.
      IRIS supports `Z/R` and `KDP/R` relationships.

3. In **CAPPI Height**, specify the height of the **CAPPI** surface in kilometers and tenths of kilometers.

    The IRIS **CAPPI** algorithm constructs **CAPPI**s by interpolating in height and range to the selected **CAPPI** surface.

    An intermediate product in cylindrical coordinates (**CAPPI** height, surface range and azimuth) is produced first, followed by the final conversion to Cartesian coordinates for the display.

    Enter a range of numbers for a 3-D **CAPPI**.

4. If needed, use the **CAPPI Fill** option.

    The interpolation algorithm requires that for each point in the output product image, there be an elevation angle both above and below the selected height.

    For example, the volume scan in the sample **CAPPI Product Configuration** menu cannot produce a **CAPPI** at 5 km (3.1 mi) height for ranges less than 5 km (3.1 mi) because there is no angle higher than 5 km (3.1 mi) in this region.

    Similarly, if you pick a low-level **CAPPI** surface, the lowest elevation angle is higher than the surface at far ranges. The resulting product display shows these unsampled areas as black.

    The **CAPPI Fill** field uses the highest elevation angle to fill the near ranges, and the lowest elevation angle to fill the far ranges, eliminating the black areas. This approach is also called **Pseudo CAPPI**. This is not recommended for 3-D **CAPPI** products.

More Information

‣ Display Parameters (page 29)
‣ Configuring Map Projections (page 21)
‣ Product Parameters (page 25)
‣ Associating Products with Tasks (page 19)
‣ CAPPI Height Selection Tool (page 255)

3.6 **FCAST**: Forecast

FCAST is an array of direction and intensity vectors used to shift a current weather display interactively.

FCAST is used in the **Quick Look Window** to project other products forward in time to see a future look at the weather. The product takes as input another image product and then determines an average speed and direction of echo motion. It is similar to the automatic **TRACK** product in that it computes the mean speed and direction by doing a cross-correlation between the current and the previous version of the product.
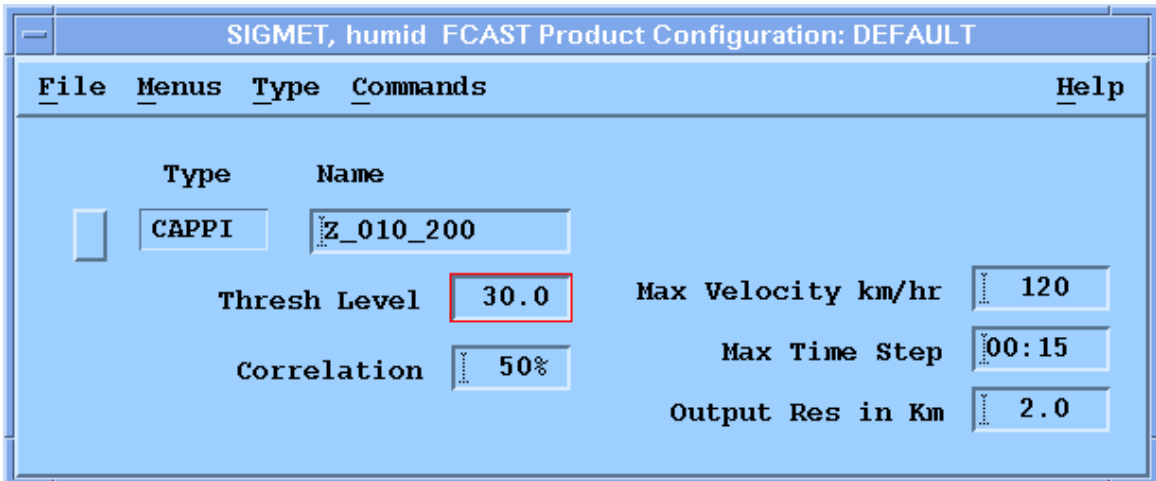
Figure 6   **FCAST** Example Configuration

1. Select **Type > FCAST**.

2. In **Type**, select the input product type and name.
   Select a product that can produce a reasonable echo motion.
   If you are using the product to forecast the motion of a **VIL** features, base the forecast on a **VIL** product.
   In the example, a **CAPPI** product of dBZ is selected.

3. Specify the **Threshold Level**.
   The algorithm applies a weighting function to the data based on this level. Values above this level are amplified and values below this level are suppressed by the weighting function.
   When you choose the level, look at a display in loop mode and get a feeling for what color boundary moves in a representative way.
   Usually you want to select a weak value to make sure it is always present when there are echoes. In the example, 30 dBZ is selected.

4. Specify the **Correlation** threshold.
   The correlation algorithm returns a speed, direction and a correlation value which serves as a quality indicator.
   The value that you enter here is used as a threshold. If the calculated correlation value in the product is larger than the threshold that you set, a forecast product is generated. Otherwise no forecast product is generated.
   The example requires a correlation of 50 %. A larger value would be more restrictive. A smaller value is more likely to pass unreliable motion estimates when the data are noisy.

5. Specify **Max Velocity in km/hr**.
   It is also possible that the forecast product has a sufficiently high correlation, but makes an absurd velocity perhaps because of random noise.
   To prevent this, enter the maximum velocity that you would expect to see for weather echoes in your area. A value of 120 km/hr (74.6 mi/hr) is a reasonable first guess.

6. In **Max Time Step**, specify the maximum time that the Forecast product looks back to find the previous image.
   The **FCAST** product calculates the mean echo motion by doing a correlation between the current and previous image.
   The example value of 15 minutes is a reasonable starting point since the average life of an individual convective cell is approximately 30 minutes.

7. In **Output Resolution**, specify the resolution of the output field of vectors.
   FCAST products produces a single motion vector. You can also make your own forecast product using your own algorithms. These can have a field of non-uniform motion vectors.
   The total size of the output field matches the size of the input product (in km). For example, if the input product has a 200 km (124.3 mi) range (400 km (248.5 mi) total span), and the output resolution is 5 km (3.1 mi), then the output is a grid of 80 × 80 motion vectors.

8. To display and use the **FCAST** product in forecasting to, select **Quick Look Window > Forecast**
   The display of the forecast product shows an array of motion vectors similar to the **WIND** or **NDOP** products.
   The **Quick Look Window** displays support additional features that are not supported by the **FCAST** product itself:
   • Non-uniform motion fields are correctly used in the forecast.
   • Non-uniform intensity changes are correctly applied.
     The forecast product creates a corresponding array of intensity changes which are all set to 1.

More Information

3.7 **HMAX**: Height of Maximum Intensity Product

Figure 7   **HMAX** Configuration Example

HMAX displays the height of the maximum reflectivity above each output pixel.

It is similar to the **MAX** product which shows the value of the maximum reflectivity over each pixel. The **HMAX** algorithm requires a volume scan. It searches vertically through the elevation angles above each output pixel and determines the height of the maximum intensity.

The product can in some cases show the presence and height of the bright band. In animation, rapidly decreasing heights in convective storms may indicate the presence of a microburst. A limitation of the **HMAX** product is described in the note below:

> **ℹ** A limitation of this product is that the heights are quantized by the number of elevation angles. It is not possible to interpolate the maximum, so the height is the height of the elevation angle where the maximum occurs (with earth curvature correction). Using a volume scan with many elevation angles and the 2D Cartesian smoother both help to mitigate the effects of this quantization.

▶ 1. Select **Type > HMAX**.

2. In **Data : Display**, select one of the following:
   - `dBT Height`
   - `dBZ Height`
   - `dBZc Height`

3. In **Max Range**, select the maximum range for the product.

4. In **dBZ Minimum**, select the smallest maximum reflectivity for which to show height.

5. Apply the **AZ/EL Smoother**.
   Because of the quantization of the heights by the elevation steps in the volume scan, you may want to try a longer smoother than you would typically use to help smooth over the quantization (for example, 3.0 km (1.9 mi)).

More Information

‣ Display Parameters (page 29)
‣ Product Parameters (page 25)
‣ Associating Products with Tasks (page 19)

3.8 **LAYER**: Layer Average

LAYER can compute layer averages of any polar data types in the ingest files.

LAYER can also convert to liquid first and compute `VIL Density`. When computing `VIL Density`, the output is in g/m**3.

The **LAYER** algorithm is similar to **VIL**, except that it divides by the layer thickness. The algorithm process is:

1. Searches out all points in the layer (accounting for earth curvature) over a given range and at a given azimuth that intercept the **PPI** scans of the volume scan, including one point above and below.
2. If the data is `Z`, it linearizes it.
   If the output is **VIL** density it converts the `Z` values to `W` values (`W` refers to water content), and averages the values in the layer.
   Each data point is assigned a weighting corresponding to the height interval that it represents in the layer.
3. The result is an intermediate **PPI** product that is transformed to a Cartesian product and stored.
4. If `Z` is selected as the product data parameter, but at run time only `T` is available (or vice versa), the product runs with the available data parameter.

LAYER can function when only one angle is in the task, but for best results, this is not recommended . If no angle in the associated task passes through the layer, no result can be calculated.

For a layer of 5 ... 10 km (3.1 ... 6.2 mi), in the volume scan example in the following figure, layer averages cannot be calculated for ranges less than 5 km (3.1 mi). In the resulting product display, there would be an inner circle of an unscanned area.

Figure 8   Example of 15-tilt Volume Scan
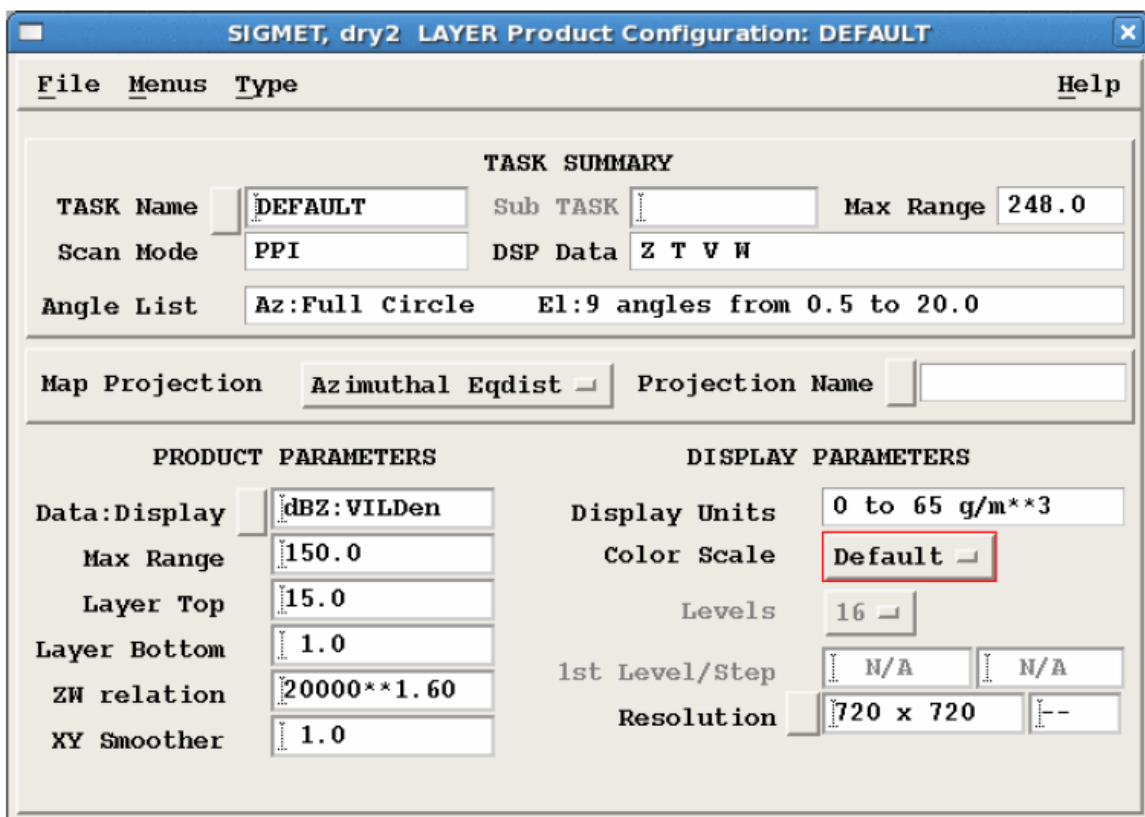
Figure 9   **LAYER** Configuration Example

▶   1. Select **Type > LAYER**.

2. In **Data:Display**, select the type of data that is computed.
   - To compute `VIL Density` data, select one of the following:
     - `dBT:VILDen`
     - `dBZ:VILDen`
     - `dBZc:VILDen`
   - To compute a layer average, select any other data type.

3. In **Layer Top** and **Layer Bottom**, select the top and bottom heights of the layer in kilometers and tenths of kilometers.

4. If you are computing `VIL Density`, in **ZW Relation**, select the reflectivity-water content (`Z-W`) relationship.
   The default value is for rain. For snow, reduce the coefficient to a smaller value, such as 2000, to account for the lower reflectivity of ice.

More Information

‣ VIL: Vertically Integrated Liquid (page 98)

3.9 **MAX**: Maximum Data with Side Panels

MAX shows the maximum reflectivity over each pixel as well as the East-West and North-South maximum projects in side panels.

The **MAX** product provides an easy-to-interpret presentation of the echo height and intensity in a single display. It is especially useful for displaying areas of severe weather.
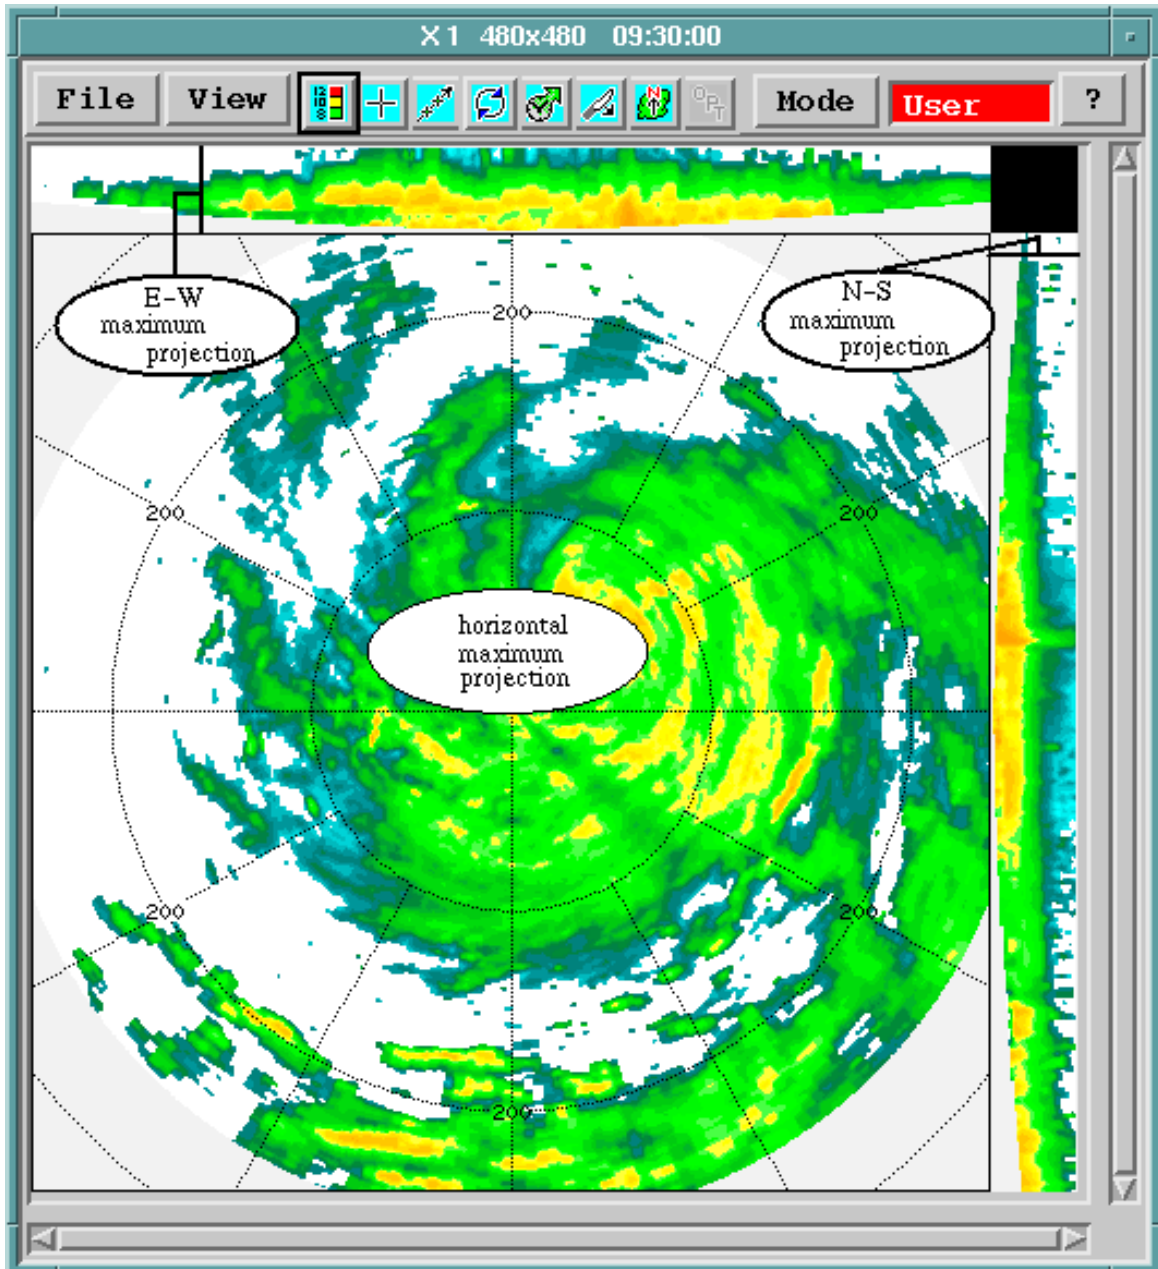
Figure 10   Example **MAX** Display

> Because the radar cannot see to the surface of the Earth, there are curved boundaries at the bottom of the side panels.

MAX is based on a volume scan task and is calculated as follows:

1. Construct a series of **CAPPI** products to span the selectable layer, spaced uniformly over the height interval. The number used is the larger of one for every 500 meters (1640 ft 5 in) or one for every 2 pixels of side panel size, up to a maximum of 50. The values are interpolated to create the side panel image.
2. Determine the maxima data values for the horizontal and vertical projections: East-West and North-South.
    *   The horizontal maximum projection is calculated from the highest data value in a user-specified layer over each pixel.
    *   The E-W maximum projection is obtained from the maximum reflectivity for each pixel along the corresponding N-S line.
    *   The N-S maximum projection is obtained from the maximum reflectivity along E-W lines.
3. Smoothing, if enabled.

3.9.1 Configuring MAX Products
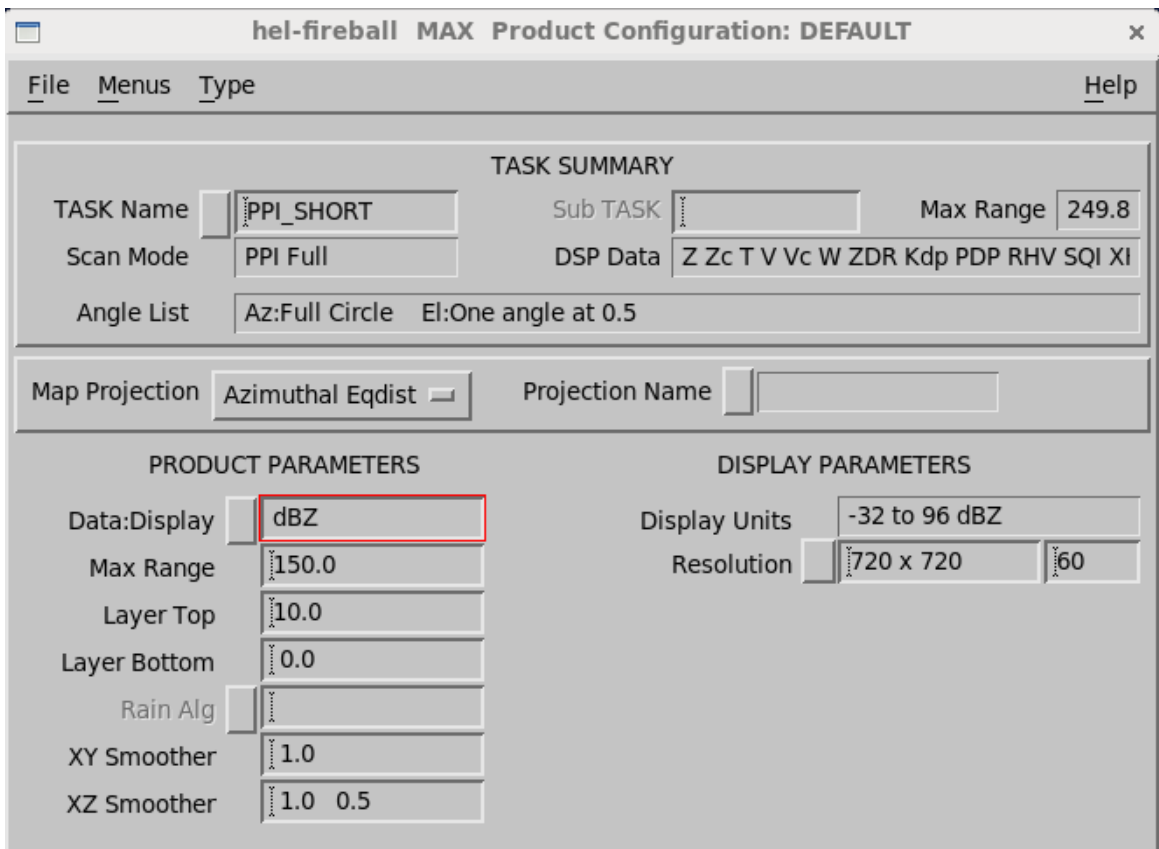
Figure 11   **MAX** Example Configuration

▶ 1. Make a diagram of the layout that you want, similar to the following figure.
       Compare this to the elevations in your task to see if you have a sensible resolution,
       especially in the upper parts of your **MAX** box.

2. Select **Type > MAX**.

3. Select a **Product Name** that is mnemonic.
   For example, `Z_0_16_220` indicates that the product covers the layer 0 … 16 km (0 … 9.9 mi) to a range of 220 km (136.7 mi).

4. Select the **TASK Name**.
   The associated **TASK Name** must correspond to a volume scan task.
   At least 2 elevation angles are required.

5. In **Data:Display**, select the input data.
   Available data types include `dBT`, `dBZ`, `dBZc`, `V`, `Vc`, `W`, `ZDR`, `ZDRc`, and `Kdp`.
   Rainrate can be output for the appropriate types.

6. Define the **Max Range**.
   The **Max Range** corresponds to the span of the image from the radar to the edge of the picture in the East-West direction.
   In the example, the range of 220 km (136.7 mi) is covered by 120 pixels to yield 2 km (1.2 mi) per pixel resolution.

7. Select the **Layer Top** and **Layer Bottom** to the nearest 100 m (328 ft 1 in).
   When selecting the layer top value, remember the maximum data height allowed for your system as configured in the **Setup** utility.
   In the example, the layer spans 0 to 16 km (0 to 9.9 mi). The corresponding side panel resolution is set to 32 pixels so that each pixel represents 0.5 km (0.3 mi). The center of the bottom pixel is at 0.25 km (0.2 mi) height and the center of the top pixel is at 15.75 km (9.8 mi) height.

8. Define the smoothing levels (in km).
   The **XY Smoother** is for the horizontal maximum projection.
   The **XZ Smoother** is for the side panels and includes values for horizontal smoothing and vertical smoothing.

   > ℹ️ In many applications, you may want to use no smoothing so the that maxima are not diminished by the smoothing filter.

9. For the **Color Scale**, **Levels**, and **1st Level/Step**, type the number of levels, the first level value, and the level step in `dBZ`.
   The color levels for the `Zmax` display are identical to those defined for the reflectivity.

10. In **Resolution**, select from the following default values or type other values in the field. **Resolution** consists of 3 numbers: the number of E-W pixels, the number of N-S pixels, and the height of the side panels.

Resolution	Pixels	Height
Low	240 × 240	20
Medium	480 × 480	40
High	720 × 720	60
XHigh	940 × 940	80

> Resolution values represent the total size of the display, including the side panels.

More Information

‣ Display Parameters (page 29)
‣ Configuring Map Projections (page 21)
‣ Product Parameters (page 25)
‣ Associating Products with Tasks (page 19)

3.10 **MLHGT**: Melting Level Height

The melting level bright band typically represents the location of frozen hydrometeors falling into air warmer than the freezing level, where the particles start to melt.

The height of the melting level is an important aspect to understanding what is happening micro-physically within precipitation systems. Changes of melting level height can assist in determining the likelihood of hail and lightning, and is related to estimating strength of downdrafts due to phase change from ice to liquid. During winter conditions determining the location of liquid versus ice precipitation at the surface is improved with knowledge of the melting level, especially in mountainous terrain.

The presence of the bright band associated with the melting level must be accounted for with quantitative interpretations of weather observations.

In IRIS and RDA software, the **SRI** product uses the melting level height to perform Vertical Reflectivity Profile corrections, to apply Fall Speed corrections to radial velocity, to distinguish between convective and non-convective precipitation, to account for the different attenuation between liquid and ice, and as a membership function in the Hydrometeor Classification function. Thus an automated determination of the melting level height and how it varies spatially and temporally adds value to the overall weather radar system software.
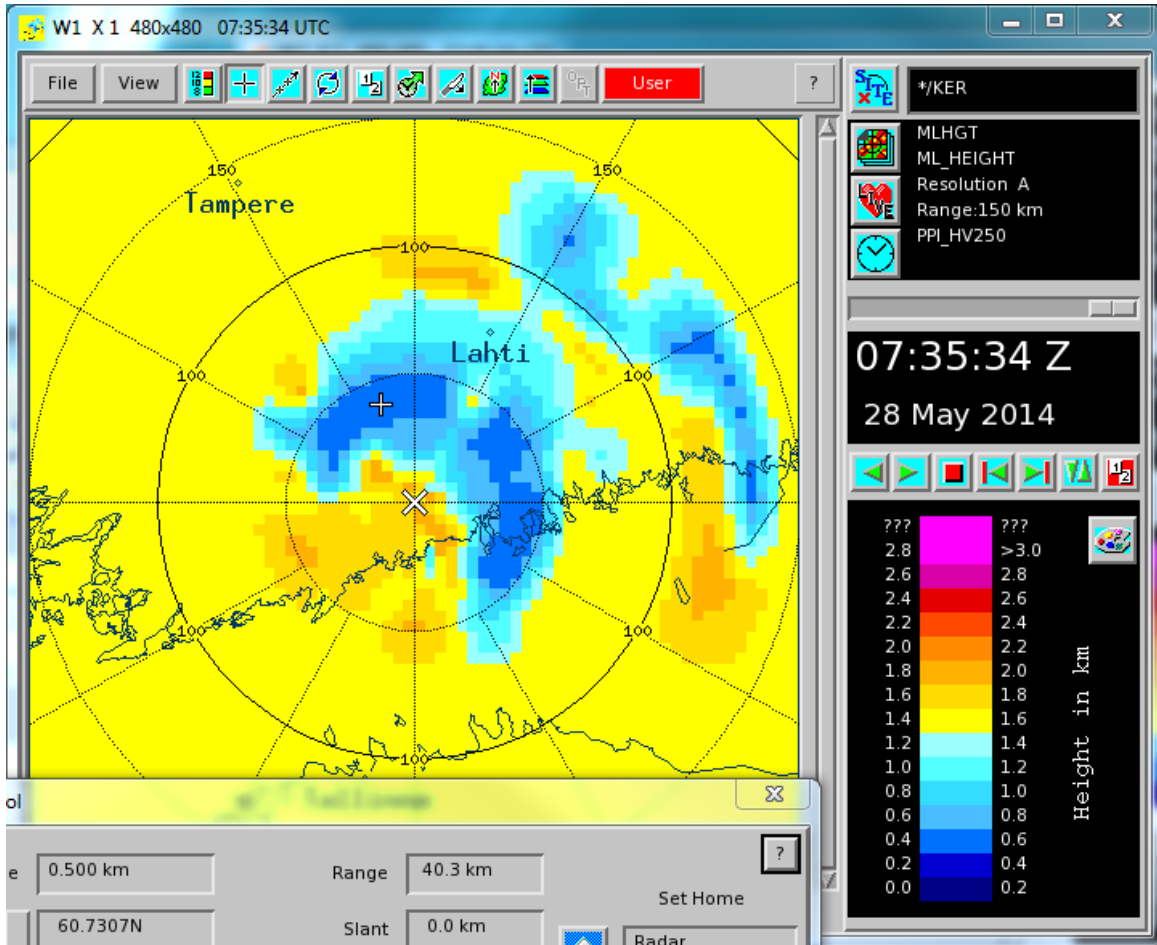
Figure 12   Melting Level Height Detection

The **MLHGT** product produced by IRIS is a depiction of the melting level heights on a Cartesian map. The melting level heights are determined on a user defined grid providing an illustration how the melting level heights changes geographically.

Dual Polarization radar data is used as input along with a priori information. As with other IRIS products, the **MLHGT** product may be animated in time and overlayed with other data to increase situational awareness of the precipitation structure in the atmosphere.

3.10.1 MLHGT Algorithm

IRIS dual polarization raw volume data is used as input for the **MLHGT** algorithm. This data may come from either single **PPI**, volume, **RHI**, or sector scans. The algorithm also functions independently from scan geometry matching data from the scan polar co-ordinate system to an Earth relative co-ordinate system. However as with any radar observation, a higher number of elevations in a volume scan provides the best results.

The following figure shows the **MLHGT** algorithm flow diagram. The top level illustrates the radar observation inputs, followed by prior information. The blocks represent consecutive functional steps. A key qualifier for each step is expressed in parenthesis.
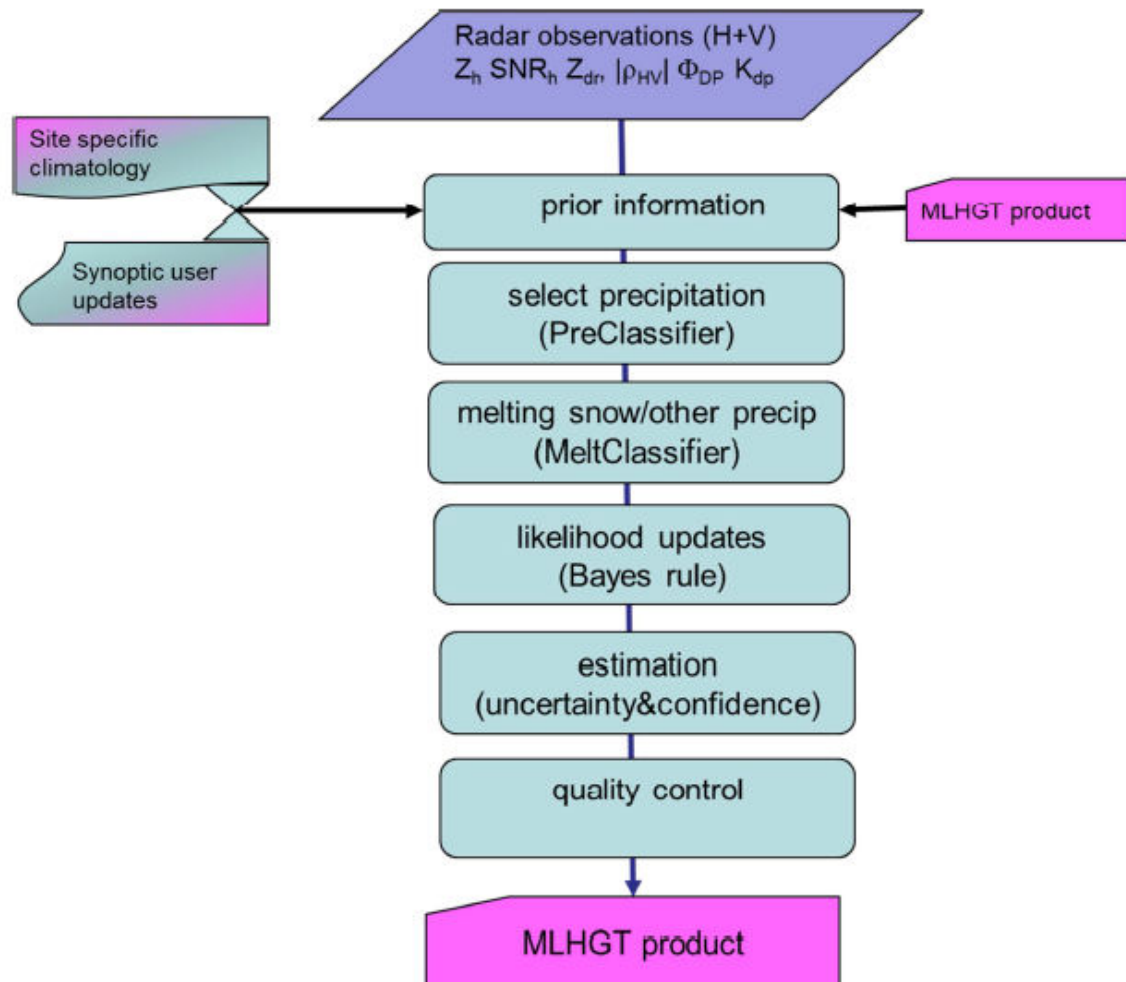
Figure 13   MLHGT Algorithm Flow Diagram

The first step in the processing is to ensure only precipitation particles are being considered for the melting level height estimates. The `PreClassifier` tags each range bin as being from precipitation or non-precipitation targets. The non-precipitation flagged bins are removed from the considered data.

The precipitating bins are then classified into 'melting snow' or 'other precipitation' by a dedicated fuzzy algorithm known as `MeltClassifier`. `MeltClassifier` is similar to the `MeteoClassifier` used in `HydroClass`. However, the `MeltClassifier` substitutes a Signal to Noise Ratio (SNR) Membership function (MBF) in place of the melting level height MBF used by the `MeteoClassifier`. The SNR MBF is configured for low SNR to be part of the 'melting snow' and high SNR in 'other precipitation'. Other particle types such as rain, dry snow, hail, and graupel which are normally characterized or also grouped into the 'other precipitation' class.

The statistical method known as Bayesian Rule is used to infer the likelihood of a region being either 'melting snow' or 'other precipitation', using a priori information and posterior conditioning. The a priori information is the first estimate of the 0 °C heights from an outside source. This could be automatically inserted from radiosonde or NWP data, or could come from the climatological values found in IRIS/Setup.

The likelihood of the **MLHGT** can be updated by each subsequent radar observation over time building up a confidence in the estimate. If there are a high number of 'melting snow' or 'other precipitation' bins in a region, and the ratio of one condition versus the other is high, confidence increases quickly. Likewise, if there is a low number of classifications in the region or if the ratio of 'melting snow' to 'other precipitation' is almost equal, the confidence remains low. Once the minimum confidence is reached the melting level height form MLHGT is assigned to that region, else the height remains to be the climatological value. The Bayesian inference can be conceptually understood as constructing a display of vertical cross-sections in which the most likely position of the ML can be visually recognized by inspection of the `MeltClassifier` decisions.

The regions are presented as columns of data in an Earth co-ordinate system. The user is allowed to configure the vertical resolution and azimuthal size of each region. The columns size in range domain are scaled to match their size in azimuth. The columns range extent grows proportionally as the size in azimuth as distance from the radar increases. Thus the columns are effectively the same shape throughout the domain. This is advantageous because the statistics is improved at farther distances, due to an increasing number of data points, which compensates for the loss in resolution by beam broadening.

3.10.2 Configuring MLHGT Products

Figure 14   **MLHGT** Configuration Example

Use the IRIS **MLHGT Product Configuration** menu to configure the **MLHGT** product.

▶   1. Select **Type > MLHGT**.

2. Configure the **MLHGT** product:

Variability

The **MLHGT** algorithm uses climatological melting level altitude as an initialization parameter in the Bayesian inference.

The **Variability** field sets the realistic limits of allowed **MLHGT** solutions centered at the climatological melting levels configured in Ingest/Setup utility.

For example, if the climatological melting level height is typically 3 km (1.9 mi) above sea level and the variability is set at 3000 m (9842 ft 6 in), the accepted range for the solution is from 1.5 ... 4.5 km (0.9 ... 2.8 mi).

> You can still use the `setup_change` function to automatically update 0 °C heights from other sources, such as radiosonde or NWP, on a more timely basis.

Hgt Bin Spacing

An Earth co-ordinate grid system is created where a **MLHGT** value is derived for each grid point. This field sets the vertical bin spacing centered on each grid point.

Data in the vertical bins are used as input to the height determination. Allowable inputs are between 50 ... 2000 m (164 ft 1 in ... 6561 ft 8 in).

Az bins

The spatial resolution varies as a function of range due to the broadening of the radials. This field defines how many column sectors exist in the 360° of view. For example, if you add a value of 24, there are 24 columns for each range interval, each having a size of 15° (360/24 =15).

> By specifying a small base scale of the grid, more `MeltClassifier` estimates are needed as input to achieve a given confidence level.
> Specifying a large grid scale lowers the overall resolution.
> You may wish to achieve something in between, which provides sufficient data availability at the desired confidence levels with higher resolutions.

Decay Time

As confidence of the **MLHGT** estimates decreases, either by lack of radar observations or latency in time, the **MLHGT** Bayesian approach smoothly transitions the estimates from the radar observed height to the climatological melting level height at every grid point.

The **Decay Time** field sets the length of the time window for this smooth transition. The default value is 1 hour.

Min Confidence

At each grid point the **MLHGT** algorithm checks the confidence in the radar **MLHGT** estimate by comparing it with those in the neighboring grids. The confidence is accepted if it exceeds the minimum confidence level defined in this filed AND a minimum number of neighbors meet the confidence. Else the confidence in the estimate is set to a very low value.

Min Cluster

Defines the minimum number of neighbors needing to achieve the confidence.

3.10.3 Using MLHGT Data Flow with Other Products

To make full use of the **MLHGT** product, it should be distributed to other IRIS nodes to support product generation and user display. The **MLHGT** product should also be returned to IRIS/Radar and the RVP900 signal processor at the radar site.

The following figure is an example of **MLHGT** product distribution. The figure assumes the RAW product is received at a centralized IRIS/Analysis node. The IRIS/Analysis node is configured to generate a **MLHGT** product. This node would have the climatological settings of the 0 °C heights. Optionally, there may be custom external process that obtains more realistic 0 °C from other sources and updates the 0 °C stored in temporary memory. The node generates the first **MLHGT** product according to the saved configuration and saves it as a file. IRIS then manages this product file in a similar to other types of products.
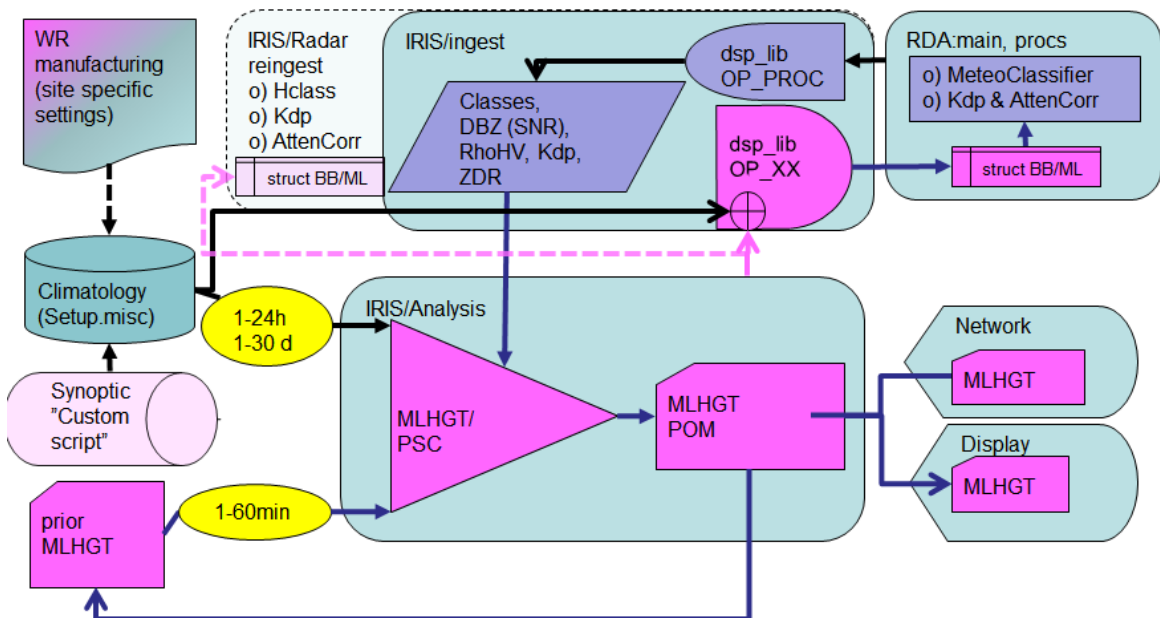
Figure 15   MLHGT Data Flow for Use with Other Products

Ideally the **MLHGT** product is disseminated from the IRIS/Analysis node to the IRIS/Radar node located at the radar site. IRIS/Radar automatically shares the information in the **MLHGT** product to the RVP900 RDA software process through `OP_CODE` in the `DSP` API. See *RVP900 Digital Receiver and Signal Processor User Guide*.

RVP900 uses the knowledge of the melting level for real-time attenuation correction, real-time hydrometeor classification, and `PhiDP` conditioning for adaptive `Kdp` computation. When the original ingest files are created, the radar derived melting level is also stored in the metadata. The ingest files containing the radar derived melting level heights are typically transferred back to the IRIS/Analysis for use in product generation.

3.11 **PPI**: Plan Position Indicator
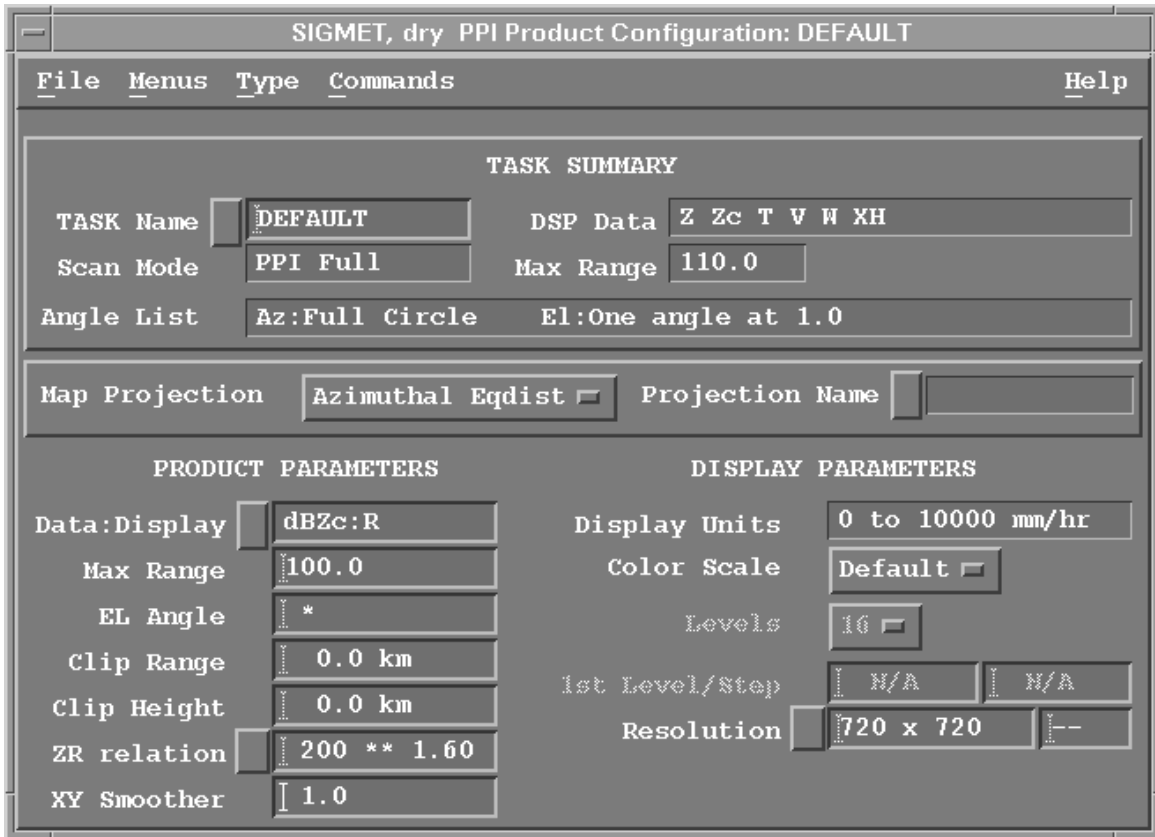
Figure 16   **PPI** Example Configuration
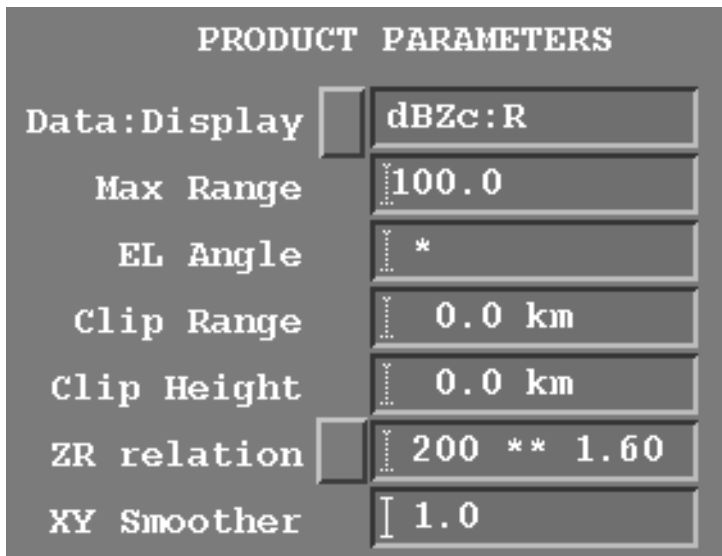
PPI is a full screen image used primarily for weather surveillance purposes.

PPI shows the distribution of the selected data parameter on a constant elevation angle surface, the classic radar display. IRIS **PPI** products are made as soon as the requested elevation sweep is completed rather than at the end of a volume scan.

A **PPI** shows data from one elevation, all azimuths. You can use it at the radar site for quality control.

1. On the IRIS menu, select **Menu > Product Configuration** .

2. Select **Type > PPI**.

3. Select the button next to **TASK name**, and select a task from the list.

4. Select the button next to **Data:Display**, and select a parameter to be displayed, typically **Z** for reflectivity.

5. Enter **Maximum Range** (km) and elevation angle which suit the task you have selected. You can see task information in upper part of the menu.

6. Select **File > Save as** and give a name for the **PPI** file.

More Information

‣ Display Parameters (page 29)
‣ Configuring Map Projections (page 21)
‣ Product Parameters (page 25)
‣ Associating Products with Tasks (page 19)

3.12 **RAIN1**: Hourly Rain Accumulation
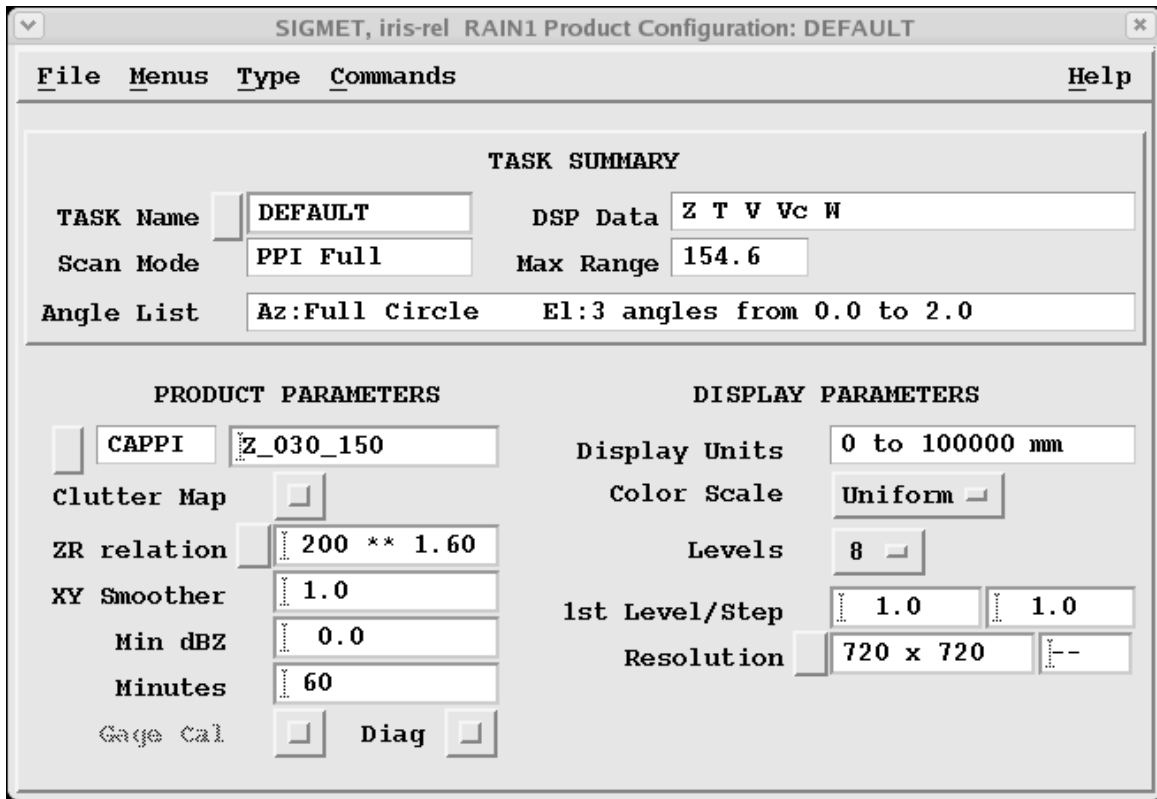
Figure 17   **RAIN1** Example Configuration

RAIN1 is hourly rainfall accumulation.

The **RAIN1** product uses **CAPPI** or **SRI** data over a defined time interval, generally 60 minutes, to obtain an estimate of the rainfall that fell in that hour. The results can be displayed or used to derive the **RAINN** or **CATCH** product.

A general-purpose volume scan task usually serves for the rainfall product. Note that the accuracy of the rainfall product degrades with range because the beam is larger at greater ranges. For example, the beam for a 1° beam width is 3 km across at 180 km range.

> If you have a Doppler radar, select Z as one of the task data outputs. This is corrected for ground clutter effects.

▶   1. Select **Type > RAIN1**.

2. Select the product type (either **CAPPI**, or **SRI**) and the name of an input product.
   The input data type can be corrected dBZ, uncorrected dBZ, or rainfall rate.
   If you are integrating rainfall rate inputs, the **ZR Relation** and **Minimum dBZ** fields are not applicable.
   The priority of the **RAIN1** product must be lower than that of the **CAPPI** and **SRI** products to ensure that all of the inputs are ready before the **RAIN1** product is made.
   Use the **setup** utility's **product** command to set the relative priorities of products.

3. To enable clutter maps, select **Clutter Map**.
   To use this, you must make a **RAIN1** product with a matching name and resolution in clear weather.
   This file is then subtracted from each hourly accumulation. The clutter map file is marked in the **Product Output Menu**.

4. In **ZR Relation**, set a Z–R relationship that is characteristic of the precipitation in your area and account for the type of precipitation, such as rain or snow.
   For information on using radar to estimate precipitation, Z–R relationships, and the advantages and limitations of the technique, see a standard textbook such as Battan or Doviak and Zrnic.

ZR **Standard Relationships**	
$Z = 200R^{1.6}$	The Marshall-Palmer relationship for rain
$Z = 2000R^{2.0}$	The Gunn and Marshall relationship for snow

5. In **Minimum dBZ**, set the minimum value of dBZ used to compute rainfall amounts.
   dBZ values that are less than this are treated as 0 rainfall.
   This prevents very weak echoes from being counted as rainfall.
   In most cases, set this to -10 dBZ, although you should choose a value appropriate to your site.
   Set the value to -33 to disable **Minimum dBZ**.

6. In **Minutes**, set the desired time period for rainfall accumulation.
   Generally 60 minutes is used to produce a one hour rainfall product, but you can use shorter intervals. The minimum accumulation time is 15 minutes.

7. If you have purchased the IRIS/Hydromet software license, select **Gage Cal** to have the product use the raingage correction algorithm.

More Information

3.12.1 Scheduling RAIN1 Products

The **Data Time** that you specify when you schedule the product are quantized to the time value in the product configuration for the nearest interval that ends the integration.

For example, if you want to schedule a 60 minute **RAIN1** product to integrate the precipitation in the hour from 11:00 to 12:00, set the time to 12:00. If the associated task is running at the end of an hour, the **Product Scheduler** waits for it to complete, so the task data can be included in the hour's integration.

Use the **Skip Time** field if you want only the accumulations for specific time intervals. You can set the skip time to any multiple of the accumulation period. For example, if you want a 3-hour integration only for the periods ending at 03:00, 06:00, 09:00, ..., set the **Skip Time** field to 03:00.

Set the **Skip Time** field to 01:00 so the **RAIN1** product runs every hour. This is recommended because the **RAIN1** product for multiple hour accumulations works best if there is a product every hour.

The **RAIN1** product must wait for arrival of all the input **CAPPI**s during the accumulation period.

The trigger is the arrival of a **CAPPI** that is in the next accumulation period.

There are 2 cases that the product scheduler must handle, that is, the next hour **CAPPI** either arrives or it does not arrive. Note that in the time-line figures below:

- The horizontal line represents the data time, the time of the start of a volume scan.
- The C represents an input **CAPPI** product.
- The vertical bar represents an the next accumulation period.

Case 1: Normal Case: Next hour CAPPI Arrives

The **RAIN1** product is run as soon as a **CAPPI** from the next accumulation interval arrives.

Case 2: Late Case: Next accumulation Interval CAPPI Does Not Arrive

The **RAIN1** product runs after the `Tail` and `Wait` times have elapsed from the file arrival time of the last input **CAPPI**.

`Tail` is the time difference between the data time of the last **CAPPI** and the next accumulation period. The `Wait` time threshold is configured in IRIS **Setup > Product**, where you can configure a different value for cases with raingage inputs. The `Wait` time threshold is tuned to account for any network delays.

More Information

3.13 **RAIN-N**: N-Hour Rain Accumulation

Figure 18   **RAIN-N** Example Configuration

RAIN-N is rainfall accumulation of the last `N` hours, where `N` is selected by the user.

RAIN-N is a product of a product: You can sum any number of hours of individual **RAIN1** products. The only limitation is the number of **RAIN1** products stored on disk. The product output shows the last `N` hours of accumulation.

1. Select **Type > RAIN-N**.

2. Specify the product name.
   When you specify the name for the **RAIN-N** product, it is a good idea to indicate the integration time in the name, in the example, `03_HOUR`.
   This makes it easier to identify the product in the **Product Scheduler** and **Product Output** menus.

3. Configure the **RAIN-N** product:
   Rain1
   > Specify the product name of the input hourly **RAIN1** product by selecting the **Rain1** button and choosing from a list of products or by typing the name in the field.
   > When you enter a product name, the associated task information is displayed in the **TASK Summary** portion of the menu.

   Hours
   > Specify the number of hours to integrate. When the product runs, it integrates data for the last `N` hours.
   > If a **RAIN1** product is missing for one of the hours in the interval, the algorithm assumes that no rain fell during that hour. The product output shows how many hours were integrated.

   Color Scale, **Levels**, and **1st Level/Step**
   > The color levels can be specified, but the resolution is fixed by the resolution of the **RAIN1** configuration.

More Information

‣ Hydromet Raingage Correction Overview (page 377)

3.13.1 Scheduling RAINN Products

The **Data Time** that you specify when you schedule the product are quantized to the time value in the product configuration for the nearest interval that ends the integration.

Use the **Skip Time** field if you want only the accumulations for specific time intervals. You can set the skip time to any multiple of the accumulation period. For example, if you want a 3-hour integration only for the periods ending at 03:00, 06:00, 09:00, …, set the **Skip Time** field to 03:00.

The scheduling algorithm for the **RAINN** product looks for the arrival of the most recent **RAIN1** product (after the "current hour") and then runs the **RAINN** algorithm.

There are 2 cases that the product scheduler must handle, that is, the **RAIN1** product arrives or does not arrive.

Note that in the time-line figures below:

• The horizontal line represents the data time of the **RAIN1** products that is always exactly on the ending hour.
• R represents an input **RAIN1** product.
• R* represents the current hour **RAIN1** product whose arrival triggers the **RAINN** product to run.
• The vertical bars represents even hours.

The examples are for the case of a 4-hour integration.

Case 1: Normal Case: Current Hour RAIN1 Product Arrives

In this case, the current hour **RAIN1** product arrives (R*). The algorithm runs on this and the prior **RAIN1** products that are available over the previous 4 hours.

Case 2: Current Hour RAIN1 Does Not Arrive, But Next Hour RAIN1 Does

In this case, the current hour **RAINN** does not arrive so the scheduler waits.

Eventually the next hour **RAINN** arrives, so the scheduler assumes that the current hour **RAINN** will never arrive and runs the product on the available **RAIN1**s, without the current hour data.

If the current hour **RAIN1** eventually arrives, it does not trigger the scheduling algorithm since the `Next Data Time` in the **Product Scheduler** has already advanced to the next hour. However, it is used in future **RAINN** processing, that is, for the next hour.

Case 3: Current Hour RAIN1 Does Not Arrive, but `Wait` Time Expires

In this case the **RAINN** product runs after the `Tail` and `Wait` times have elapsed.

The `Tail` time is the time on the local computer between the arrival of the most recent RAIN1 and the current hour.

The `Wait` time threshold is configured in IRIS **Setup > Product** to account for network delay.

If the current hour **RAIN1** eventually arrives, it does not trigger the scheduling algorithm since the `Next Data Time` in the **Product Scheduler** has already advanced to the next hour. However, it is used in future **RAINN** processing, that is, for the next hour.

More Information

3.14 **RAW**: Raw Data
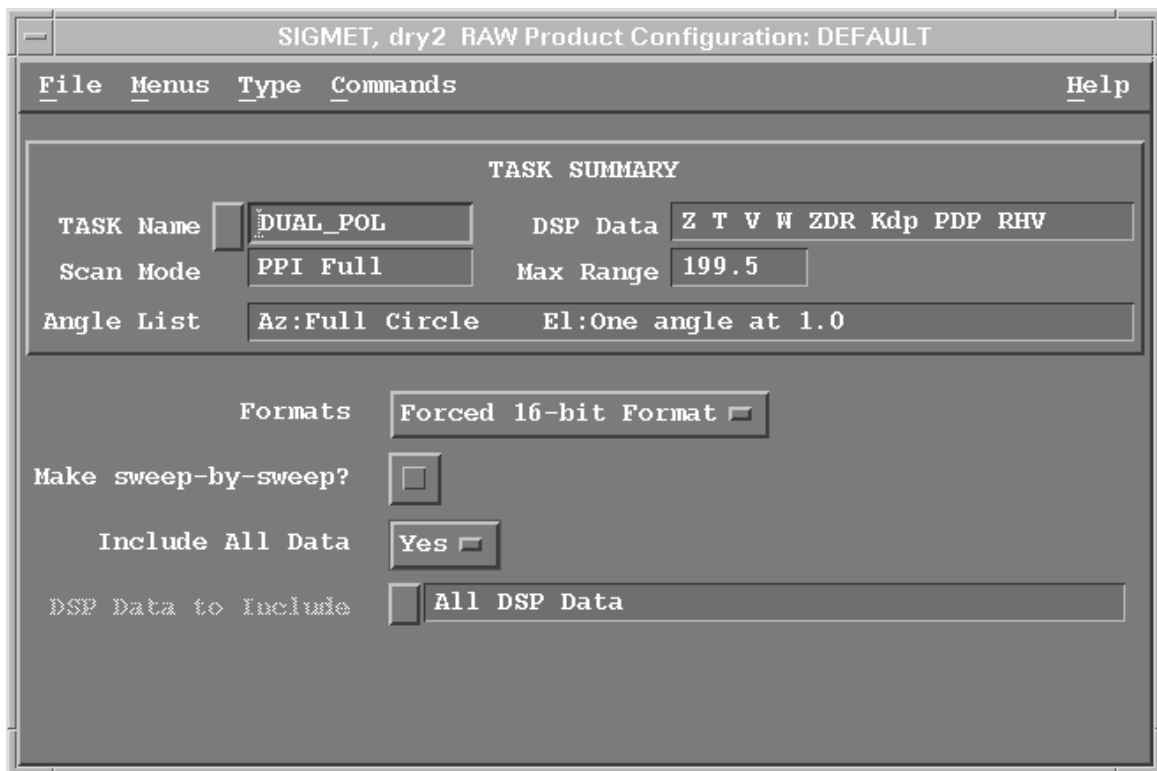
Figure 19   **RAW** Example Configuration

A **RAW** data file that contains the raw signal processor output parameters (Z, ZT, V, W, ZDR, and so on) in polar coordinates.

RAW products are made one-for-one from ingest files.

A *RAW* file contains all the information collected during a measurement task. It is compressed so it can be sent over the network to IRIS Analysis.

While **RAW** data cannot be displayed, it is useful for recording on tape/disk/DVD or transferring to another IRIS network computer for product generation.

When a **RAW** product file is either retrieved from tape or received from another IRIS computer, an ingest file is created so that products can be generated. The **RAW** product is used in the following ways:

- Offline Analysis System: **RAW** products are stored on tape and later read back to an IRIS Analysis system (not necessarily the radar host) so products can be generated for analysis or testing.
- Real-Time Network System: **RAW** products are sent to another IRIS computer and the product generation is done there rather than by the radar host. This off-loads some of the processing burden from the radar host onto other network computers that may be less busy. This is valuable if you want to generate products that are not normally produced by the radar host.

More Information

‣ Associating Products with Tasks (page 19)

3.14.1 Configuring RAW products

1. Select **Menu > Product Configuration**.

2. Select **Type > RAW**.

3. Select the button next to the task name, and select a task from list.

> Vaisala recommends you give a **RAW** product the same name as its associated task.

4. To copy the product files over a slow network connection, select **Make sweep-by-sweep? > On**

5. To filter the **RAW** product to only include some of the data types of the original data, deselect **Include All Data**.
   For example, you might want to send all data to a research customer but skip sending the dual-pol data to a simple display computer.

> Vaisala recommends that you leave **Include All Data** enabled until a need arises to limit data types.

> If you use a wild card, you cannot merge data from hybrid tasks into one file.

6. Select **File > Save as** and give a name for the *RAW* file.

3.14.2 Recording RAW Data Automatically

You can configure a **RAW** product automatically when the associated task runs.

> Because **RAW** products require substantial disk space, it is best to record them "on the fly," so they are not automatically deleted by the disk `Watchdog`.

1. Configure the **RAW** product.
   See 3.14.1 Configuring RAW products (page 67).

2. Load the associated tasks in the **TASK Scheduler** and select **Go**.

3. In the **Product Scheduler** menu, load the product and choose **All**.
   This causes the **RAW** product to be generated every time the task runs.

4. In the **Product Output** menu, specify the desired archive media as the output device and choose **All** in the header line for the product, so all output is sent to the archive.

5. In the **Archive** menu, mount a tape or disk and select **Commands > Record**.

3.15 **RHI**: Range Height Indicator
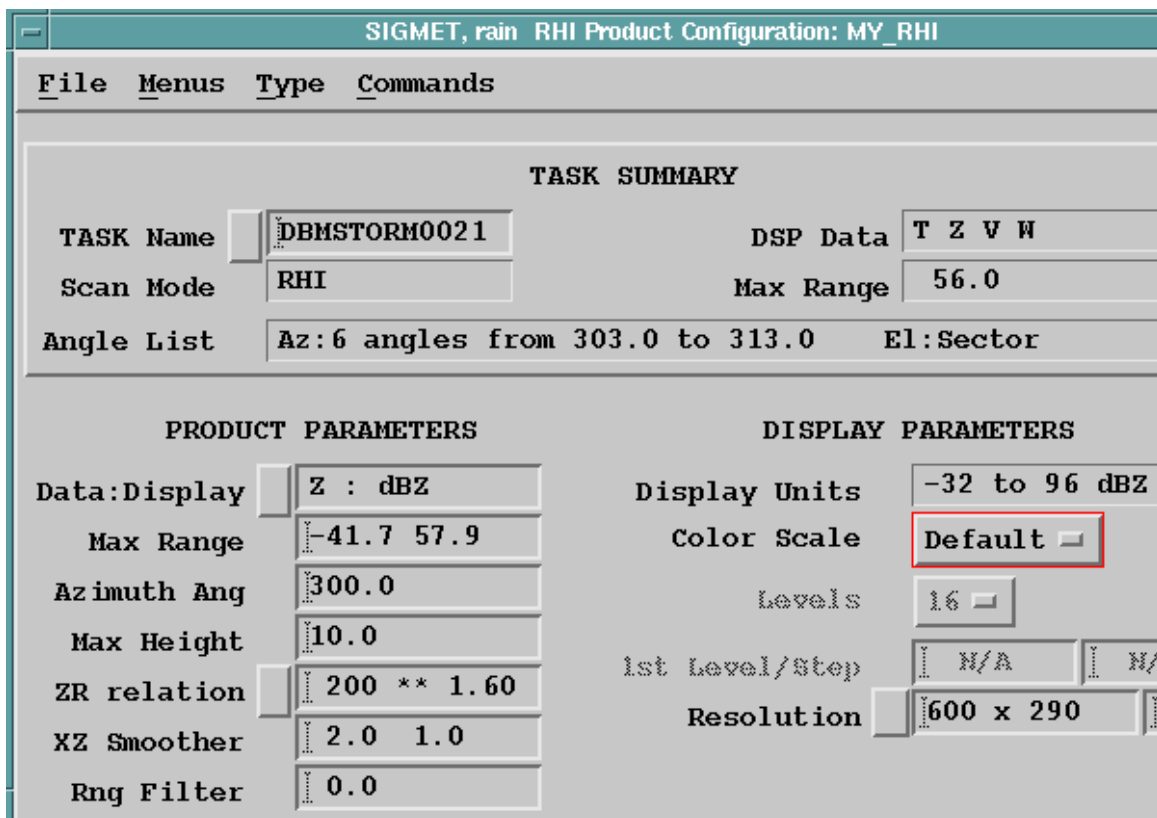
Figure 20 **RHI** Example Configuration

RHI is a full screen image showing the detailed cross-sectional structure of a storm, used for identifying severe storms, hail and bright band.

In general, you schedule the associated **RHI** task through a region of interest.

During **RHI** scanning, the antenna azimuth is fixed and the elevation is swept, typically from near 0 … 90° to create a vertical cross-section effect.

If your antenna permits, it is often useful to make an over-the-top **RHI** task (elevations for example -2 … 182). This gives a picture covering 2 elevations at the opposite sides of the antenna. Negative elevations are most useful in **RHI** if you select the azimuth to the direction of least beam blocking of the radar horizon.

▶ 1. Select **Type > RHI**.

2. In **Data : Display**, select any data type defined in your **RHI** task.
   Also, if you measure **V** you can select **SHEAR** or output data type for radial shear. Radial shear is the difference between the velocity of successive bins in range.

> **SHEAR** is allowed as an output data type to the **RHI**, but not to other similar products like **PPI**.

3. In **Max Range**, enter up to 2 values in format of `xx.x` kilometers.
   If you enter only one value, the first one is assumed to be 0.
   The first limit is the beginning of your display.
   Zero means the image starts from the radar site.
   If you enter negative values, data is picked from another azimuth on the other side of the radar site in the same task, or if your task is scanning over-the-top.
   You can also enter positive values, especially if you do not use high elevations and you want to avoid the flat corner of a triangular-shaped image.

4. In **Azimuth Angle**, specify an azimuth angle:
   If you do not know this value in advance, IRIS selects the closest angle at run time, so you are always assured of getting a display.
   If the associated task has only one azimuth angle sweep, that angle is used for the **RHI** product.
   If you enter * to the field, the product is made for every azimuth in the associated task.

5. In **Max Height**, **XZ Smoother**, and **Range filter**, specify the **Max Height** field in `XX.X` km.
   Because you specify both the maximum range and height, you can make an **RHI** at any horizontal-to-vertical aspect ratio.
   Note that the final display is from 0 to the **Max Height** above the reference height.
   The radar is located on the left edge, above the corner if the radar is above the reference. Negative elevation angles are included in that case. Rays are inserted straight, and the grid of constant height lines is curved to show earth curvature.
   You can type the **X-Z** (height) **Smoother** values in km.
   If you have **SHEAR** as an output data type, it is somewhat noisy, and must be smoothed with a range filter. This number is unused for data types other than **SHEAR**.

6. In **Resolution**, select **Low** (288 × 136), **Medium** (600 × 290), **High** (840 × 530), or **XHigh** (1060 × 750) resolution or type a resolution in the field.
   RHI products look best with a rectangular resolution.

More Information

3.15.1 RHI Cumulonimbus Example

For detailed vertical structure, there is no substitute for a good **RHI**. The following example shows a cumulonimbus seen with **RHI** from the southern part of the US in summer.

The **RHI** shows that this is a powerful, actively growing **Cb** with top at 13 km (8.1 mi). The echo over-hang and weak echo region on the right indicates a well-organized strong inflow and vertical motion, both characteristic of a severe thunderstorm. The reflectivity core (red is >56 dBZ) indicates a probable hail shaft. An upper level anvil, presumably from another **Cb** can be seen close to the radar.

Figure 21   **RHI** Cumulonimbus Example

3.15.2 Should I Use RHI or XSECT?

You can study the three-dimensional structure of the atmosphere with the **XSECT** and **RHI** products, which show a range-height representation or vertical "slice" of a radar parameter.

If you have areas of interest, such as paths to runways at an airport, you can run **RHI** or **XSECT** products associated to those in you routine schedule.

Feature	RHI	XSECT
Radar location	Radar location is fixed to one point.	You can make a cross section at any point along the polar volume. Note that you cannot get below the lowest elevation.
Resolution	Typically, **RHI** gives a much better resolution.	--
Schedule	If the user of the radar images can control the tasks, they can define a product configuration in advance (using * for the azimuth). When something interesting appears, they can run an **RHI** task with an azimuth determined using other IRIS products.	You can schedule an **XSECT** when you want, even ad hoc. If your radars are in a network of several users, or if your task schedule is tight, you perhaps cannot let the users run **RHI**s when they want.

More Information

3.16 **RTI**: Range Time Indicator
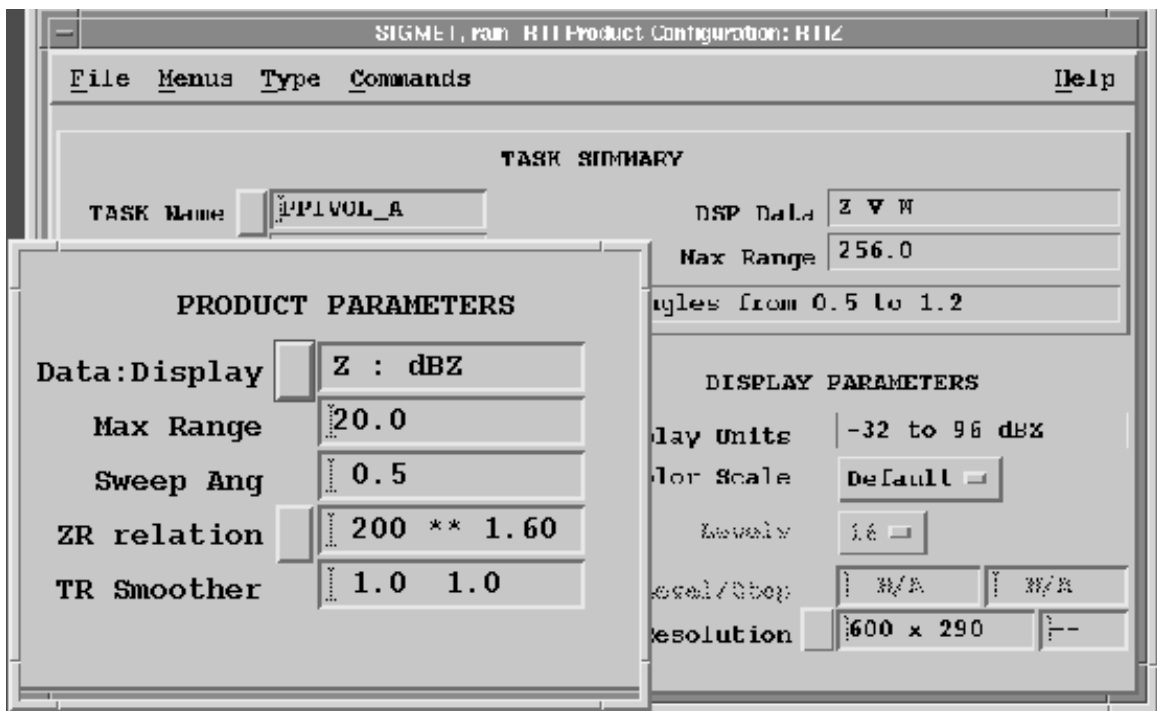
Figure 22   **RTI** Example Configuration

RTI displays time along the horizontal axis and the vertical axis displays range from the radar.

An **RTI** can be used for any scan type (for example, **PPI**, **RHI**), but is most useful for manual scans or "searchlight" scans which are manual scans at a fixed position. Perhaps the best application is for vertical incidence scanning in which case the **RTI** product shows a time-height cross-section of the atmosphere.

The following example shows an **RTI** and a **PPI** from the same situation. From both pictures, the maximum range is 50 km (31.1 mi) and the color scale is same. In this case, **RTI** has been made from a **PPI** task rotating 10° per second so the horizontal axis is both seconds from start and azimuth in tens of degrees. You can see the weak echo region near 220°, starting from 15 km (9.3 mi) in both. This can be beam blocking, which is probably easier to recognize in the **RTI**.

1. Select **Type > RTI**.

2. In **Data:Display**, choose any parameter available in the task configuration.

3. In **Max Range**, type the maximum range in kilometers.

4. In **Sweep Angle**, select any of the elevation angles in the task.
   The angels are listed in the upper part of the menu.

5. If you have selected rainfall rate (R) for output, in **ZR relation** define how to transform reflectivity (dBZ) to rainfall intensity (mm/h).

6. In **TR Smoother**, define the smoother in XY directions of the **RTI** image.
Typically, a small value is used for research purposes and a bigger value for customer products.
The time direction smoother is in seconds, the range direction in km.

7. In **Resolution**, choose from low, medium, high, and extra high.
The default resolutions are rectangular to match the default rectangular window sizes of the **Quick Look Window**.

More Information

‣ Display Parameters (page 29)
‣ Product Parameters (page 25)
‣ Associating Products with Tasks (page 19)

3.17 **SRI**: Surface Rainfall Intensity

SRI (Surface Rainfall Intensity) provides input for the **RAIN1** product to obtain the best possible estimates of accumulated precipitation even at longer ranges from the radar.

Vertical reflectivity profile is the most important source of error in radar rainfall measurements in cool and moderate climates.

The upper parts of precipitating clouds typically give a weaker echo than the cloud base, except near melting layer where the echo is much stronger. Thus a correction is needed to estimate surface rainfall intensity.

IRIS **SRI** allows you to apply your local knowledge by providing ways to input information about the reflectivity profile, as well as methods to make educated first guesses.

SRI distinguishes convective cases from large scale precipitation, and applies the correction only to the latter. For convective precipitation, the value of the lowest clutter free bin is shown.

3.17.1 SRI Algorithm

The **SRI** reflectivity profile and melting level estimation are not perfect, but they can improve the rainfall estimates compared to performing no correction. The typical corrections obtained are of the order -10 … +5 dBZ (in mm/h scale up to factor of 4) depending on the melting level altitude, distance from radar, and the lowest elevation angle.
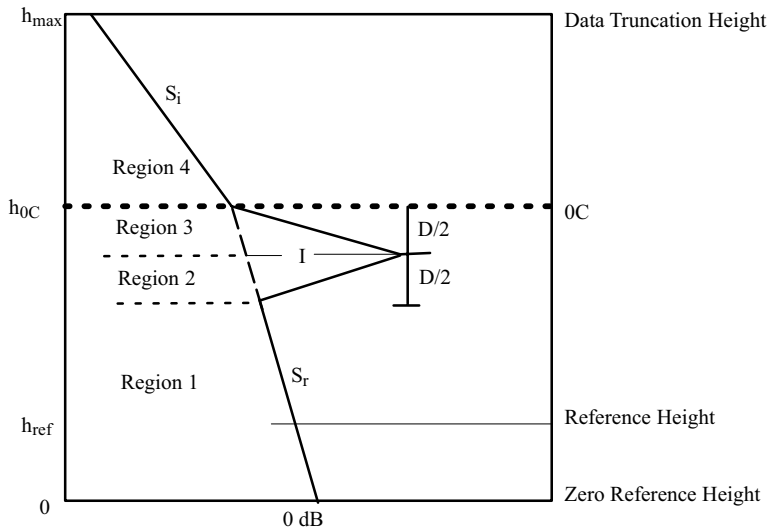
Figure 23 **SRI** Reflectivity Profile Example

The profile includes the following:

- Height `h=0` corresponds to the **Setup > Product** `Zero Reference Height`.
  This is typically set at either the nominal ground level or sea level.
- The max height corresponds to the **Setup > Ingest** `Data Truncation Height`.
- The reflectivity varies linearly in dBZ above and below the bright band.
  Separate slopes are used above and below the bright band (S_r for rain and S_i for ice).
- The bright band starts at the 0°C level ($h=h_{0C}$), has depth (`D`) and intensity (`I`), defined as the intensity difference between the peak and the intensity the rain would have at the center of the bright band, that is, determined by the continuation of the rain slope into the bright band.

The surface rainfall intensity at each pixel is calculated by finding the lowest clutter-free bin, and bringing the measured reflectivity there down to reference level by making 2 corrections:

1. The correction for the beam weighted averaging.
2. The adjustment for the profile to obtain the reflectivity at the reference height.

The **SRI** product also supports using a terrain map to determine the height that that radar beam is corrected to. See 3.17.6 Format of Terrain Map File (page 86).

Convective Check Algorithm

In convective precipitation there is usually no detectable bright band or perhaps a very weak one. This is thought to be caused by the types of particles near the melting level that are typical of convective precipitation. These are usually heavily rimed snowflakes, graupel, frozen drops (carried aloft), or hail. Since these particles tend to fall more rapidly than snow aggregates, they do not contribute to a bright band in the same way as large wet snow flakes falling at 1 m/s. That is, there is no convergence of large wet particles.

In cases where the precipitation is convective, it is not appropriate to perform a profile correction that includes a bright band and serious errors could result from doing this. In the **SRI** product this is handled by detecting convective regions and not applying any profile correction. That is, the measured value of the lowest clutter free bin is assigned to the reference height. This is done because in convective regions, the vertical reflectivity in the lower altitudes tends to be rather constant in height.

Each range bin is checked to determine if it is a convective range bin. The approach is to run an echo **TOPS** algorithm in the **SRI** product for a selectable threshold. Note, this is done in the **SRI** so there is no need to define a separate **TOPS** algorithm.

The product, in cylindrical coordinates, is made at the same resolution as that selected for the **SRI** product. It is then used as a mask for determining convective regions. That is, any bin in a ray for which the top height is more than a selectable height above the melting level is assumed to be convective.

3.17.2 Input Reflectivity Profiles

You can define the reflectivity profile with the following parameters:

- Reflectivity gradient above the melting layer (7 dBZ/km (11.3 dBZ/mi))
- Reflectivity gradient below the melting layer (1 dBZ/km (1.6 dBZ/mi))
- Melting layer height (in tenths of kilometers)
- Melting layer thickness (1 km (0.6 mi))
- Melting layer peak intensity (7 dBZ)

The melting level is attached to the ingest file. In a network of several radars and a separate IRIS Analysis machine, the profile information is given to the IRIS instance at each radar site.

> The melting layer height is given in IRIS **Setup > Product**.
> You can change the value without restarting IRIS in the **Setup** utility. This means you can adjust the reflectivity profile frequently based on external data sources such as numerical model or temperature soundings.

Getting Reflectivity Profile Input Data

The most accurate way to determine the melting layer height is to study measurement from the same radar at the same time as when the profile correction is needed.

You can see the melting layer nicely in IRIS **VVP** and **XSECT** products of either reflectivity or, even better, vertical velocity. However, this is not always available: when the precipitation area is approaching the radar, there is data only from the upper parts of the profile.

Figure 24   **VVP** Reflectivity Profile Example

Numerical weather prediction models and balloon soundings are good sources of temperature profiles.

A bulk method is to take average temperature (daily or even monthly averages) and calculate the melting level height assuming the moist adiabatic lapse rate 6.5 °C / km. This lapse rate is a good estimate during the precipitation event.

If nothing else is available, start with the following table.

Table 12   First Guess Monthly Melting Height Values for the Northern Hemisphere

Month	Polar climate	Mid latitudes	Tropics
January	-1	1	2.5
February	-1	1	2.5
March	-1	1.5	3

Month	Polar climate	Mid latitudes	Tropics
April	0.2	2	3.5
May	1.3	2.5	4
June	2.2	3	4.3
June	2.5	3.5	4.5
August	2.0	3.5	4.5
September	1.5	3	4.3
October	0.5	2.5	4
November	-1	1.5	3.5
December	-1	1.0	2.5

What if There Is No Melting Level ?

When it's snowing, there is no melting layer (-1 in the table above). Then the profile becomes a simple descending line, defined by parameter set.

- Reflectivity gradient above the melting layer 7 dBZ/km (11.3 dBZ/mi)
- Reflectivity gradient below the melting layer 0 dBZ/km (0 dBZ/mi)
- Melting layer height 0 km (0 mi)
- Melting layer thickness 0 km (0 mi)
- Melting layer peak intensity 0 dBZ

The most demanding task is to get the profile right when the melting layer is close to surface. Even small errors misplace the bright band and thus lead to severe overestimation and underestimation close to the radar. If you cannot change the profile frequently, and the temperature fluctuates below and above 0 with bright band appearing and disappearing, Vaisala recommends you apply the snow profile as described above. You still have the bright band overestimation problems, but at least you fix all the weakening above the bright band.

> Make sure you give the altitude information from external data source referring to the same reference height (sea level, antenna level) as defined in the IRIS setup. Be careful when the bright band is close to ground.

3.17.3 Setting the Current Melting Level Without Restarting IRIS

Normally, when you change the **Setup** files, you have to quit and start IRIS to have the changes valid.

For **SRI** this is not practical, since in the case of fast moving front you might want to update the melting layer height every hour.

To make changes to setups while IRIS is running, you must push information through a pipe to a program called **setup_change**.

1. Open a terminal window and type the following command:

```
$ echo "iris_setup.misc.ifallspd_melts[1]=20" | setup_change -load
```

This changes the February melting level to 2.0 km (1.2 mi). Note the number of the month is UNIX style table index, thus 0 is January, 1 is February and similar. The units are in 1/10 of km. You can change many of the profile and melting level setups on the fly.

2. To make several changes at the same time, put the data in a file as follows:

```
$ echo "iris_setup.misc.ifallspd_melts[1]=20" >testfile
$ echo "iris_setup.misc.ifallspd_melts[2]=25" >>testfile
$ cat testfile | setup_change -load
```

Alternatively, put several **name=value** pairs in the same string if separated with **"\n"**. You can **ssh** this string from a central computer to all your various radars. This interface is primarily intended to be invoked from automated script files. To see the current values type the following:

```
$ setup_change -list | grep ifallspd
```

3. To specify that the melting level is unknown, you set this value to valid in **Setup**:

```
"iris_setup.misc.lMeltHeightValid = 1"
```

If you stop providing this information, you must set this value to invalid.

3.17.4 Configuring SRI
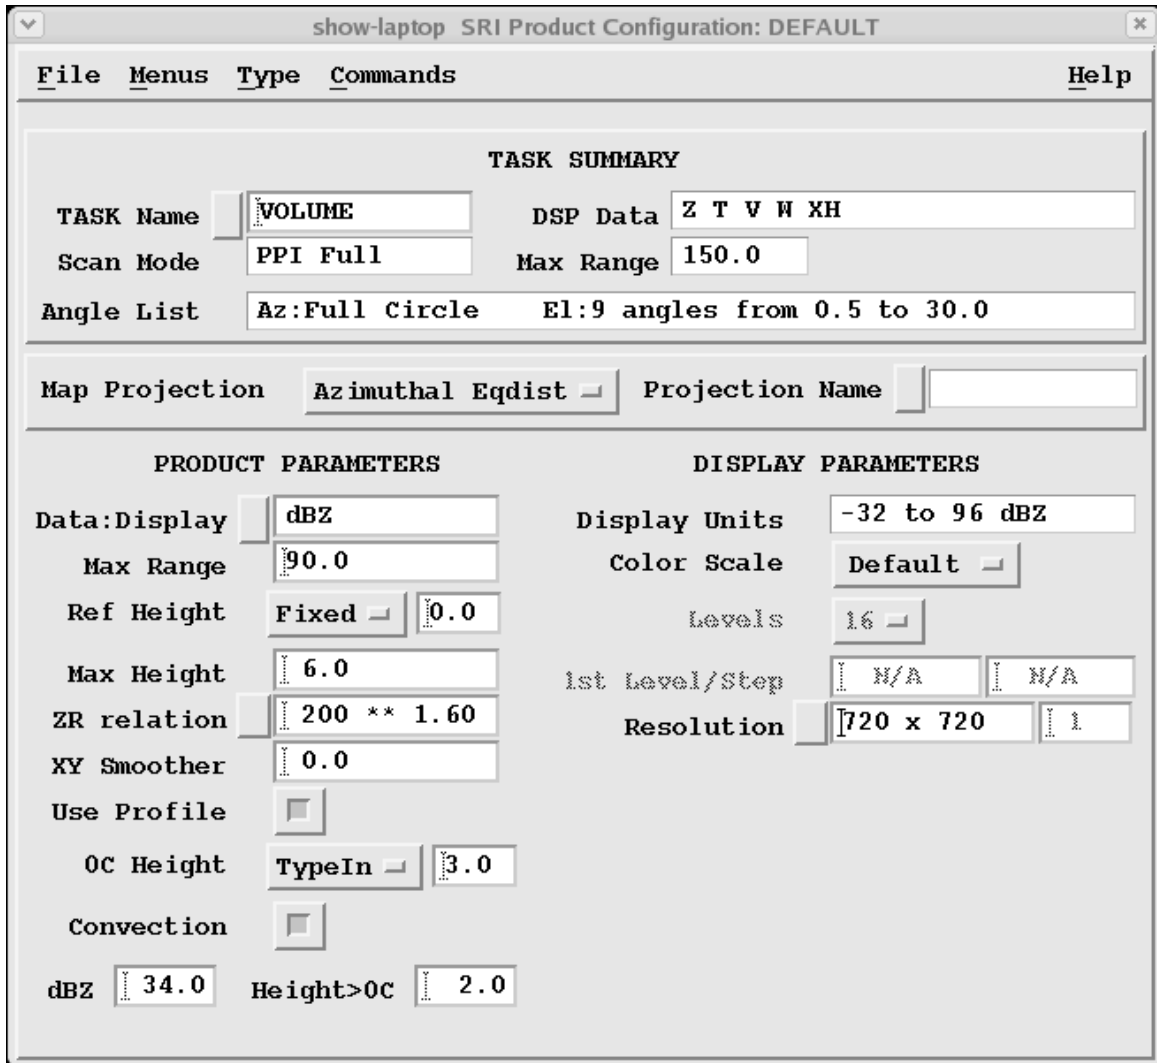
SRI (Surface Rainfall Intensity) provides input for the **RAIN1** product to obtain the best possible estimates of accumulated precipitation even at longer ranges from the radar.

▶   1. Select **Type > SRI**.

2. For **DSP Data** and **Data:Display**, **SRI** products require T and Z values as input.
   Make sure that both are recorded in your task.
   For **Data:Display** you can select if you want to see the result in dBZ or in intensity mm/h (dBZ:R) .
   In dBZc and dBZc:R, the c refers to the included corrections for attenuation, beam blockage, and similar as configured for the task. Even if you display mm/h, the convection threshold is given in dBZ.
   For radar systems that import data, the **SRI** product can also run without T (Total power) data. In a such a mode, the "lowest clutter-free bin" feature is unavailable, but the product can still be run.

3. In **Ref Hgt**, select the reference height to which the reflectivity profile of the **SRI** product is calculated.
   - **Fixed**
     Type the height (km). This is height above 0 reference.
   - **Map**
     IRIS retrieves the height information from a terrain map file. See 3.17.6 Format of Terrain Map File (page 86).

4. In **Max Height**, set the maximum height where the "lowest clutter free bin" is searched for.
   If there are no clutter free bins below this height, no correction is performed.
   You must set a reasonable level, such as below 5 km (3.1 mi). Note that at far ranges, when the lowest beam is above the **Max Height**, the lowest beam is used.

5. To use a profile in the correction, enable **Use profile**.

> This is an easy way to switch the correction temporarily off. You do not need to change other places of your production chain.

6. In **OC Height**, select the height of the melting level with one of the following options:
   - **Ingest**
     The height of the melting level is extracted from the **Setup** on the radar computer at the time of the measurement. (This is important if you process archived data.)
     All the other gradient numbers are pulled from the current setup on the product computer.
     If you have a network of radars and make the products centrally in one IRIS analysis, you can set an individual melting level for each radar at the **Setup** of the IRIS radar and use the **Ingest** setting here.
   - **Setup**
     IRIS uses the melting level heights given in **Setup > General Setup** on the product computer.
     If you want to use the same melting level for all radars and keep adjusting just one value, use the **Setup** option here.
   - **Type-In**
     When you are testing the product interactively, use **Type-In** to type the value.

7. Convection is an area of strong (precipitation) echo above melting layer. For
   Convection, do one of the following:
   - Enable **Convection** to be able to define criteria to determine whether precipitation is
     convective or large scale, and only the latter is corrected.
     - In the first box, type the threshold (in dBZ). Determine this value using the
       intensity of convection in your locale.
       Typically a value of 35 … 40 dBZ is appropriate for most locations in mid
       latitudes, 20 dBZ in cool climates.
     - In the second box, type the convective top height (how much above melting layer
       it should be, in km) in the next box.
       This is typically 1 … 2 km (0.6 … 1.2 mi) to allow some clearance above the bright
       band.
       The suitable value for threshold height depends of the accuracy of the melting
       level height estimate. The smaller the value, the more often pixels are considered
       convective.
   - Disable **Convection** to have all precipitation corrected with the above given profile.

More Information

‣ Configuring Map Projections (page 21)
‣ Display Parameters (page 29)
‣ Product Parameters (page 25)
‣ Associating Products with Tasks (page 19)

3.17.5 SRI Profile Correction Example

This example shows the profile correction in a typical situation with bright band at 3 km
(1.9 mi), gradient below the bright band (in rain) 3 dBZ/km (4.8 dBZ/mi) and above the
bright band (in snow) 10 dBZ/km (16.1 dBZ/mi). Radar is located at height of 0 km (0 mi)
and the reference height is at 1 km (0.6 mi). Full power half beam width is 1°.

The correction is calculated for 2 elevations, 1.0° (in black) and 2.0° (in grey). Note that the
correction is negative near the radar where measurement it made below the reference
height and that the curves have different shapes since the bright band is filling the Gaussian
beam only partially.
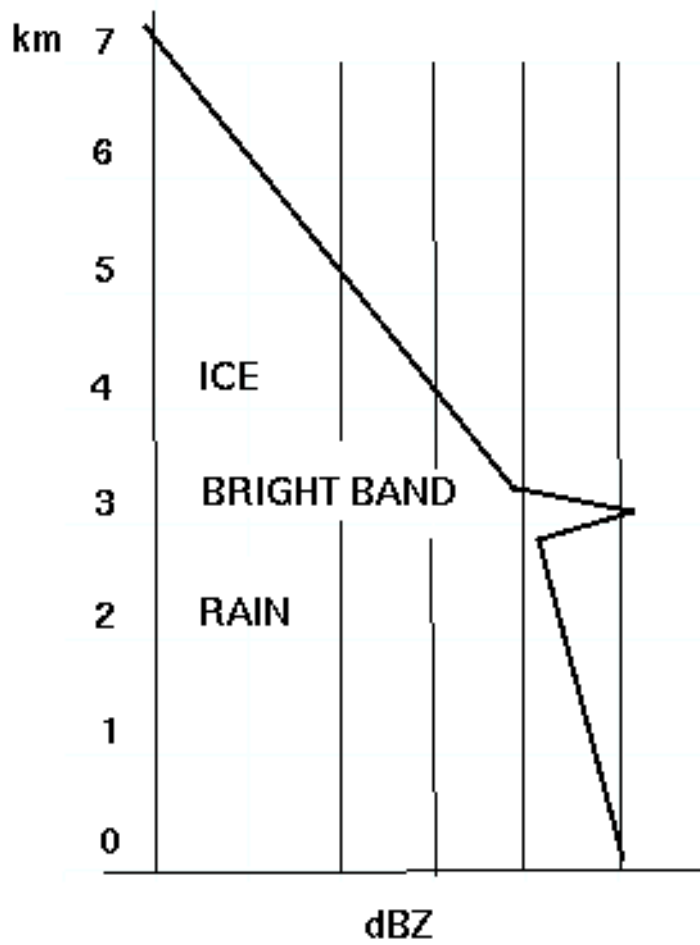
The reflectivity profile in the case of the example. Note, that at high altitudes the profile soon reaches values below the minimum detectable signal MDS, and or the cloud top.

Figure 26   Example Profile Correction against Range

Profile correction in dB as function of distance from radar for elevation 1.0° (black) and 2.0° (grey).

3.17.6 Format of Terrain Map File

The terrain maps are stored in the *surface_height.conf* file in the *${IRIS_CONFIG}* directory.

This file contains information for each radar site stored in an ASCII polar format. Here are the first few lines of such a file:

```
SITE Boston
AZ_RESOLUTION 1.0
RANGE_RESOLUTION 1.0 300.0
AZ 0
5
3
3
3
8
14
15
12
10
```

1. The **SITE** command indicates the site name for data until the next **SITE** command.
2. The **AZ_RESOLUTION** specifies the azimuth resolution of the file. In this case it is 1°. We expect to find 360 different azimuths in the file. Match this to your normal azimuth resolution, though it finds the closest in any case.
3. The **RANGE_RESOLUTION** specifies the range step in km, and the total number of km we are going to.
4. In this case, 1 km (0.6 mi) steps out to 300 km (186.4 mi).
5. Next is the **AZ** command, followed by the surface altitude in meters for each km in range.
6. After 300 altitudes there is another **AZ** command and then possibly another **SITE** command.

3.17.7 SRI References

- *Koistinen, J. (1991). Operational Correction of Radar Rainfall Errors due to the Vertical Reflectivity Profile. In Proceedings of the 25th International Conference on Radar Meteorology, AMS 1991, p.91–94*
- *Joss, J. and Pittini, A. (1991): Real–time Estimation of Vertical Profiles of Radar Reflectivity to improve the Measurement of Precipitation in an Alpine Region. In Proceedings of the 25th International Conference on Radar Meteorology, AMS 1991, p. 828–831*
- *Vignal, B., Galli G., Joss J. and Germann, U. (2000): Three Methods to Determine Profiles of Reflectivity from Volumetric Radar Data to Correct Precipitation Estimate Journal of Applied Meteorology 39(10) 1715–1726*

3.18 **STAT**: IRIS System Status

A **STAT** file contains a report indicating the state of all key components of the radar software.

STAT (status) products are typically made on all IRIS systems that are critical to system operation, such as the radar host, the product generator, and or critical display systems. Instead of containing radar data, status products contain information about how well the IRIS system is working.

These products are shown in the product output menu, and can be transferred to other IRIS systems using the product output menu.

When transferred to another IRIS system, the receiving system becomes aware of faults or problems with the sending system. For example, the radar-host status products are distributed to most systems in a radar network, and the red-X is displayed on various windows based on these status products.

The status product lists faults if any of the following occurs:

- BITEX critical faults
- RCP communication failure (RCP **DEAD**)
- RVP signal processor error
- IRIS internal critical fault. These are internal errors that are flagged as critical in IRIS.

> **i** If there is a change in status such as a critical fault, the **Status** product is generated immediately.

▶ 1. In IRIS **Setup > Product**, enable status product generation and configure the product parameters.

Status product generation
If enabled, IRIS also generates status products when the overall system status changes from OK to Fault, or the other way around.

Time between status products
Define how often you want to generate status products (in minutes). Select 0 if you do not want to generate these products.

Make product for each task
Instructs IRIS to make a new status product each time a task starts.
This is useful if you wish to do status product task slaving.

Status Prod maximum file count
Type the number of status product files to be kept on the system.

STATUS product receive timeout
Select **Disabled** or **Enabled**.
If enabled, set this time slightly longer than the time between status products.
The default value is 11 minutes with products mode every 10 minutes.
This allows IRIS to detect failures of other networked systems. If it has been more than the specified time since a status product has arrived from another system, it is marked as `Timeout` on the network status display.

2. Identify critical systems.

3. Identify separate receiving system(s).

4. To configure the network identifiers, select IRIS **Setup > RCP > Network Status Reports to RCP**.
This makes it possible to identify the reported sites.

5. Output status products to receiving systems.

6. Determine which status products are referenced by the *SIGNALS.DAT* file for action.

7. Configure the *SIGNALS.DAT* file to take the appropriate action.

8. In the **Product Output** menu, configure the status products for output.

The status products now enable IRIS systems to alert other systems on the network that there are problems or that a particular event has occurred.

These messages can be parsed by the *SIGNALS.DAT* file and actions can be taken such as sending an email to a cell phone.

3.19 **THICK**: Echo Thickness Product

Figure 27   **THICK** Example Configuration

THICK shows the thickness of cloud echoes.

THICK is the same as the difference between the **TOPS** and **BASE** values.

When calculating **THICK**, for each output pixel in the product, the algorithm searches downward through successive elevation angles to find the height at which the signal goes above the user-specified **dBZ Contour**. Similarly it searches up from the lowest elevation angle to find the height at which the signal goes back down below the contour. The thickness is the difference between these values.

The final output of the product is a color-coded map of echo thickness heights for the selected **dBZ Contour**. You can also use **IRIS Quick Look Window > OPT** to display the average reflectivity in the same layer.

THICK is affected by the same limitations as the **TOPS** and **BASE** products.

1. Select **Type > THICK**.

2. In **Data : Display**, choose `dBT Height`, `dBZ Height`, or `dBZc Height`.

3. In **Max Range**, choose the maximum range for the product.
   Note that earth curvature effects are greater at greater ranges. You can display a **PPI** at your lowest elevation angle and use the cursor tool to see what heights are possible to observe at various ranges.
   For example, for a 0.5° elevation angle at 100 km (62.1 mi), the beam is centered at 2.3 km (1.4 mi). Therefore it is not possible to detect echo bases lower than this.

4. In **dBZ Contour**, select the `dBZ` threshold for the echo base.

5. Select **Pseudo** for all pixels at which the lowest elevation angle is above the contour to use the height of the lowest elevation angle as the base.
   Similarly, all pixels at which the highest elevation angle is above the contour use the height of the highest elevation angle as the top.

6. Apply the **AZ/EL Smoother**.
   Enter the values in degrees, first for the azimuth and then for the elevation direction. Typical values are 1.0, 1.0.

More Information

‣ Display Parameters (page 29)
‣ Product Parameters (page 25)
‣ Associating Products with Tasks (page 19)

3.20 **TOPS**: Echo Tops

Figure 28   **TOPS** Example Configuration

TOPS is a color-coded contour map of the top of a selected dBZ level.

Either **Z** or **ZT** can be used as the basis for the estimate.

Heights are displayed in kilometers and tenths of kilometers. A **PPI** volume scan task is required (either **FULL** or **SECTOR**).

The **TOPS** product is an excellent indicator of severe weather and hail. For example, a 50 dBZ top 1 km (0.6 mi) above the freezing level can be produced only by a vigorous convective storm, and is most probably caused by the presence of hail. For air traffic applications, the search can be made using a lower threshold value, such as 10 dBZ, to determine the height of surrounding precipitation.

1. Select **Type > TOPS**.

2. Configure the **TOPS** product:
   Data : Display
   - Z Height
   - T Height

   The product can bederived from either Z or T data. If Z is selected as the `Product Data` parameter, but at run time only T is available (or vice versa), the product runs with the available data parameter.

   dBZ Contour
   You can select the threshold contour level in dBZ. The **TOPS** algorithm then makes a downward search at constant range in cylindrical coordinates to determine when the threshold is crossed. It then interpolates in height to obtain the height of the threshold contour. Similar to the **CAPPI**, there must be an elevation tilt both above and below the contour to obtain a valid top height.

3. When the downward search detects a dBZ value greater than the threshold, there are the following cases:
   Normal Case
   The dBZ value in the next higher ray (elevation angle) at the same surface range is used to interpolate the height of the contour that must lie in between.

   Undetected Echo Aloft
   If there is no detectable echo in the next higher ray aloft, then an interpolation cannot be done. However, a top must exist somewhere between these 2 rays. In this case, the top height is taken as the height of the point where the detection of dBZ >> Threshold was made.

   No Ray Aloft
   If there is no higher elevation angle aloft, then a top exists somewhere aloft, but there is no information as to how much higher the top may be. In this case a special color (indicated by "**???**" in the legend) denotes that an indeterminate top exists.
   For example, at ranges less than 5 km (3.1 mi) in the volume scan shown in the following figure, the algorithm cannot determine tops that are above the highest elevation angle.

Figure 29   Example of 15-tilt Volume Scan

More Information

3.21 **TRACK**: Track/Forecast

Figure 30   **TRACK** Example Configuration

TRACK products show the projected motion for storm features (centroids) based on a series of input products from different times.

The **TRACK** product consists of a series of track points connected by lines, or "tracks." Each track shows the motion of one centroid over a specified time span, plus a forecast point showing where the centroid gives its current direction and speed. A **TRACK** product may contain multiple tracks if more than one centroid shows up in the input product, as shown in the following figure.
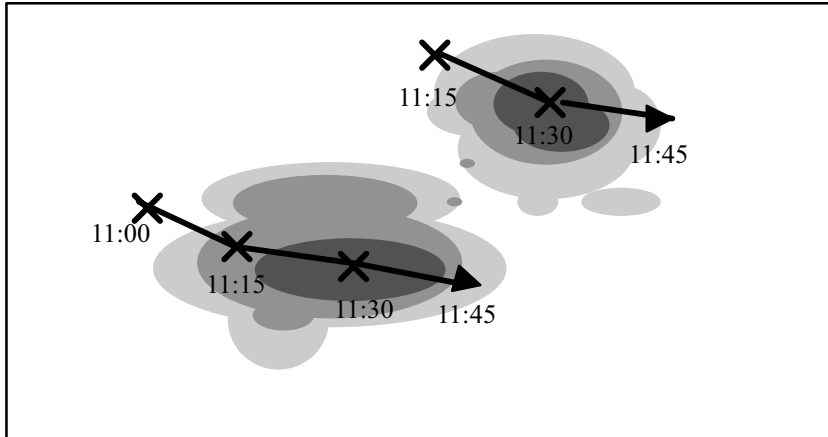
Figure 31   **TRACK** with Two Centroids

The input product can be any horizontal Cartesian product, such as a **PPI** product. The **TRACK** product applies the same logic as the warning product to locate the centroids within the input product. You define the threshold level and size of the centroids, below which weather features are ignored. When new data comes in, the **TRACK** product compares the previous **TRACK** product to the new data to obtain a motion vector. This information is used to define the new **TRACK** product, including the forecast point. Warnings are issued if a centroid hits, or is forecast to pass through, a protected area.

As new data comes in, the **TRACK** product must determine whether a new track point is an extension of a previous track, the start of a new track (when a centroid is born), or the end of a track (when a centroid dies). Track parameters, such as the maximum velocity and tolerance for new points, influence this determination.

TRACK is an overlay product created in the **Quick Look Window**.

1. To create a new **TRACK** product, select **Type > TRACK**.

2. To load an existing **TRACK** product, select **File > Open**.

3. Configure the **TRACK** product:

Product Type and **Product Name**

Select the input product type and then the product name that corresponds to the selected type. Note that there must be a product configuration file for the type of product you pick.

In the case of products received over the network, you can make a product configuration file or copy the file over the network.

Centroid Threshold Level

The **TRACK** product considers only those centroids whose values meet or exceed this threshold level.

The units of measure depend on the selected product. For example, a **TOPS** threshold is specified in km, while a **VIL** threshold is specified in mm.

> If you are uncertain about the units of measure, check the appropriate **Product Configuration** menu.

Threshold Area in Sq Km

Enter the minimum size of a centroid region (in sq km). Areas that do not meet or exceed this size are not tracked.

For example, for a 3 km by 3 km (1.9 × 1.9 mi) size area, enter "9".

Max Time Span

Enter the maximum time allowed between the newest and oldest track points to be included in the **TRACK** product.

When a new input product is processed, points that are older than this time span are removed.

You set the time using the plus and minus button to increase and decrease the hours or minutes. Select **Ok** to exit the window. The time you specify is inserted into the field.

You may also type a time value directly in the **Time window** and select **Ok**.

Max Time Step

Enter the amount of time that can elapse between two points in the same track. If more than this amount of time has elapsed between the last point in a track and a new point coming in, the new point is considered the beginning of a new track.

Max Forecast Time

Enter the amount of time to project ahead for forecast points.

The forecast time should be comparable to the lifetime of the feature that you are tracking.

For example, for an isolated severe storm, a forecast time of 30 minutes would be typical. For a squall line, a period of 1 hour could be used.

You need some experience with the types of weather in your part of the world.

Max Velocity

Enter the maximum velocity allowed between two points in the same track.

If this velocity would be exceeded by a new point coming in, that point is considered the beginning of a new track.

A typical value for this field is 100 km/hr (62.1 mi/hr).

Tolerance for New Points

Enter the margin of error that can exist between a new point and a forecast point. If the new point is within the tolerance limit, it is added to the end of the track and a new forecast point is created.

If it falls outside the limit, a new track is begun.

Protected Areas for Warning Alert

Specify one or more protected areas.

A warning is generated if a centroid falls within or is forecast to move into a protected area.

Select **Protected Areas** to see a list of areas. You can toggle choices on or off. When you exit the list, the names of the selected areas are displayed in the field. Protected areas are configured with the **Setup** utility. They are constrained to be rectangles with arbitrary size and orientation angle.

You can configure up to 32 areas. (If you make a change to **Setup**, you must restart IRIS for the change to take effect.)

TDWR Style

IRIS supports two formats of warning messages. In TDWR format, only the strongest centroid in the highest priority area is reported along with its strength. For example: `"MBA 3MF 30K-`, in spoken language: ″`Microburst Alert, Three Mile Final, 30 Knot loss`″.

The older IRIS format reports all centroids in all protected areas, for example: `3 MBA warnings at 11:30 in: 3MF, 2MF`. These messages are displayed at the bottom of the display screen and optionally spoken.

Say Warning

Select this option to have the product generator to verbally signal the warning message in addition to displaying it as a message.

In the IRIS menu, you must also select **Setup > Enable Speech**.

Make Diagnostic Results

In addition to the **TRACK** product, you can create a thresholded version of the input product, useful for testing purposes. Unless you encounter a problem with the **TRACK** product, you can turn this feature off.

3.22 **VAD**: Velocity Azimuth Display

Figure 32   **VAD** Example Configuration

VAD is a display of the mean Doppler velocity at a given range as a function of the azimuth angle as the radar antenna rotates through an azimuth scan at a constant elevation.

Min-Max Rng

Select the range over which to average the velocities.

The **VAD** product averages all velocities in this range for all rays, and store a structure for each ray. This structure includes the average value, standard deviation, data count, as well as the azimuth, elevation, and sweep number.

Unfolding

Select this option to enable unfolding the data before generating the product.

More Information

‣ Associating Products with Tasks (page 19)
‣ VAD Output Options (page 256)

3.23 **VIL**: Vertically Integrated Liquid
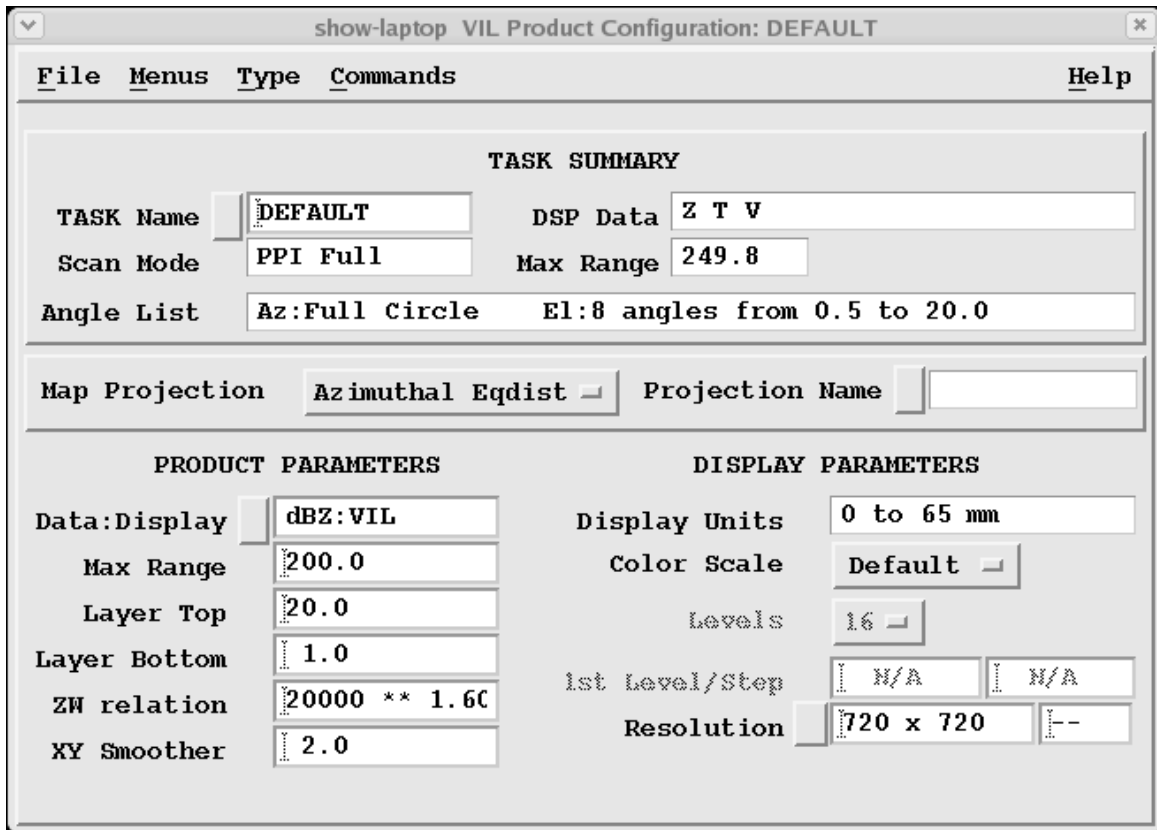
Figure 33   **VIL** Example Configuration

VIL is a color-coded map of the estimated depth of water (in mm) contained in a selectable atmospheric layer. This is an excellent indicator of severe storms.

Either **Z** or **ZT** can be used as the basis for the estimate.

VIL can compute several different values over an altitude interval or layer in the atmosphere, including integrated liquid, or averaged reflectivity. These can be excellent indicators of severe storm activity, especially with regard to the rainfall potential of a storm.

Because **VIL** can be set to look over the entire depth of the atmosphere, it is good at seeing precipitation aloft that is not reaching the ground, which can be missed by a **PPI** or **CAPPI**.

If the layer height is above the freezing level, high **VIL** values are an excellent indicator of severe storms and hail. If the layer height extends from the surface up to 3 km (1.9 mi), then the **VIL** values serve as a forecasting guide as to how much precipitation is likely to fall during the next few minutes.

When computing integrated liquid data (**VIL**-data), the output shows the estimated precipitation (in millimeters) contained within the user-defined layer. This number is sometimes also labeled in kg/m**2.

1. The **VIL** algorithm first searches out all points in the layer (accounting for earth curvature) over a given range and at a given azimuth that intercept the **PPI** scans of the volume scan, including one point above and below.
2. The algorithm converts the `Z` or `T` values to `W` (water content) values and integrates the values in the layer.
   Each data point is assigned a weighting corresponding to the height interval that it represents in the layer. The result is an intermediate **PPI** product that has the total water content as a function of surface range and azimuth.
3. Finally, the intermediate product is transformed to Cartesian and stored. If `Z` is selected as the **Product Data** parameter, but at run time only `T` is available (or the other way around), the product runs with the available data parameter.

When computing Layer Average Reflectivity (LAR-data), the output is stored in normal reflectivity. The processing is nearly identical to **VIL**-data, except that the `dBZ` inputs are converted to linear `Z` instead of `W`, and we divide by the layer thickness in the end. The average `Z` is then converted back to dB.

1. Select **Type > VIL**.

2. In **Data:Display**, select which type of data is computed.
   - `dBT:VIL` Select input type to compute **VIL** data.
   - `dBZ:VIL`
   - `dBZc:VIL`
   - `dBT` Select input and output type to compute the layer average.
   - `dBZ`
   - `dBZc`

3. In **Layer Top** and **Layer Bottom**, select the top and bottom heights of the **VIL** layer in kilometers and tenths of kilometers.

   > ⚠️ **CAUTION!** The bright band biases the **VIL** measurements. Select the **VIL** layer to avoid the freezing level height.

4. In **ZW Relation**, select the reflectivity-water content (`Z-W`) relationship.
    This field is desensitized if we are computing a reflectivity based product rather than a
    water based product. A default value for rain is provided. For snow, reduce the
    coefficient to a smaller value, such as 2000, to account for the lower reflectivity of ice.
    VIL can function when only one angle is in the task, but this is not recommended for
    best results. If no angle in the associated task passes through the **VIL** layer, no **VIL** can
    be calculated.
    For a **VIL** layer of 5 to 10 km (3.1 to 6.2 mi), in the volume scan example in the following
    figure, **VIL** cannot be calculated for ranges less than 5 km (3.1 mi). In the resulting
    product display, black would be displayed in this region to indicate that IRIS did not
    even sample the region.

Figure 34   Example of 15-tilt Volume Scan

More Information

‣ LAYER: Layer Average (page 44)
‣ Display Parameters (page 29)
‣ Product Parameters (page 25)
‣ Associating Products with Tasks (page 19)

3.24 **VVP**: Velocity Volume Processing

Figure 35   **VVP** Example Configuration

VVP (Velocity Volume Processing) provides line graphs or time against height cross-sections of wind speed, wind direction and divergence against height.

VVP obtains the following parameters averaged over a volume centered about the radar:

- Horizontal wind speed and direction
- Particle vertical velocity (airspeed and particle fallspeed combined)
- Horizontal divergence (related to vertical air motions)
- Horizontal deformation and axis of dilatation (related to frontal forcing)
- Average reflectivity
- Average **RhoHV**

You can select which data is calculated and store the resulting product. You can also turn data and be on or off. Note that for the top 4 velocity-related buttons, any button turned on implies those above are also turned on. You must also turn on at least one data type.

This algorithm is similar to the **VAD** technique, with an improved analytical approach. [1]

---

1)   *See Waldteufel and Corbin (1979, Journal of Applied Meteorology, p. 532).*

A Doppler radar can only measure the component of wind either towards or away from the radar. This is called the "radial wind." However, by looking at the wind over all azimuths around a full circle, the average wind speed and direction can be determined. This is shown schematically in the following figure for the case of a South wind. The figure shows the radial velocity at a range bin as a function of azimuth. When the radar points directly south, the radial velocity measures the full speed of the wind (toward the radar). The IRIS convention is that positive radial velocities are away from the radar. When the beam points East or West, there is no component of the wind toward the radar, so the radial velocity is 0. The divergence over the radar is estimated by taking the average of the radial velocity. Other kinematic parameters can be estimated, such as the deformation and the axis of dilatation of the wind.
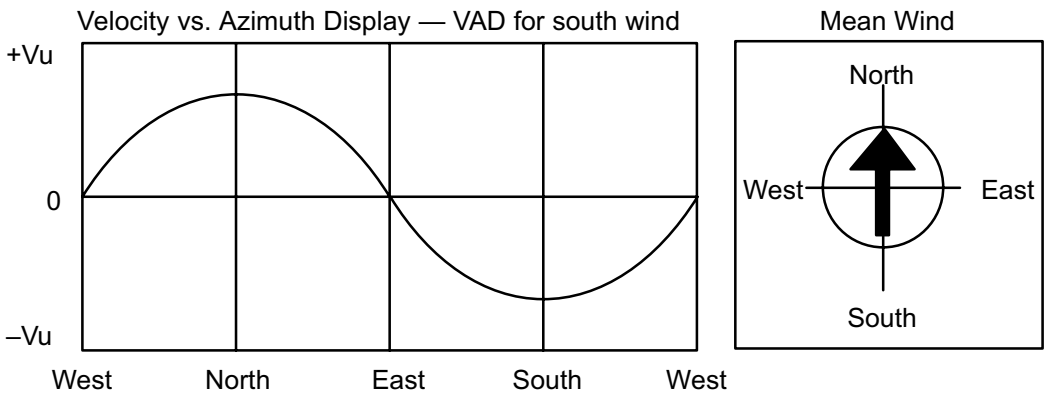
Figure 36   Example of Radial Velocity against Azimuth Display

The **VVP** algorithm assumes that the wind field varies linearly and computes the mean wind speed and direction and divergence as a function of height. For each height interval, the **VVP** algorithm performs a least squares fit of the observed radial velocity to linear windfield model. Many thousands of points are used. The output is a line graph display of each of the 3 parameters against height with error bars to indicate the uncertainty in the least squares fit.

The **VVP Product Configuration** menu requires no display configuration for the line graphs.

▶   1. Select **Type > VVP**.

2. Configure the **VVP** product:

    Min/Max Range

    Select the minimum range to avoid strong ground clutter near the radar, for example, 2 … 5 km (1.2 … 3.1 mi).
    Select the maximum range so that the highest elevation angle extends past the maximum height that you select.
    No comma is required to separate the numbers.

    Min/Max Height

    Configure the minimum and maximum height to cover the span of altitude where you want wind information.

    Height Levels

    Specify the number of height intervals. The example in the figure covers the height span from 0.5 … 15.0 km (0.3 … 9.3 mi) in 20 intervals, so one value for each 500 meters (0.3 mi).

    Min Velocity

    Specify the minimum velocity to be included in the wind volume fit.
    The recommended value is 0.5 m/s (1 ft 8 in / s).
    This is a simple way to remove bias toward 0 at low heights due to clutter which passes the clutter filter.
    A value of 0 disables **Min Velocity**.

    Bin Quota

    The maximum number of bins that is included in the least squares fit for each height interval. The maximum allowed is 10,000 bins. Vaisala recommends using the maximum of 10,000.
    If you only select the mean velocity (that is, vertical velocity, divergence and deformation are not selected), you can reduce the bin quota for faster product generation.

    Unfolding

    Unfolding can be turned on or off.
    The unfolding technique [2] , is based on selected prototype wind profiles to determine which yields the best result in the least squares fit.
    During processing, the **VVP** is run several times, and the best result is selected.
    In general, the **VVP** technique and the unfolding require at least 90° of echo coverage for reliable results.

    Reflectivity, **Vertical Wind**, **Divergence**, **Deformation**

    VVP can optionally compute the vertical velocity, deformation, and axis of dilatation of the wind field and the mean reflectivity profile above the radar.
    Select the parameters you want to compute and record.
    Note that the vertical velocity is the particle velocity. The effect of fallspeed is included.

---

2) *Siggia and Holmes, 1991 25th Conference on Radar Meteorology, Paris*

3. Configure the associated **VVP** task.

   The **VVP** algorithm performs unfolding of the data. For best results, configure a task with a large unambiguous velocity range.

   That is, use a high PRF (such as 1200 Hz) and dual-PRF velocity unfolding (either 4:3 or 3:2) so that velocity folding does not occur.

   If your standard volume scan does not use dual PRF unfolding, you might want to set up a separate task for the **VVP**.

   Four elevation angles, such as 30, 20, 10, and 5°, are usually adequate because each angles contribute to many heights.

More Information

‣ Associating Products with Tasks (page 19)

3.24.1 Producing VVP Product Output

VVP product output consists of a data file containing the **VVP** parameters for each height. The displays are computed on-the-fly from the data file and are customizable. The available **VVP** displays in IRIS are:

- Line graphs against height
  Up to 5 parameters can be selected for plotting against height. Standard deviation bars are included on the plots.
- Time-height cross-sections
  Wind barbs, wind strings or axis of dilation/deformation lines are superimposed on a colored background of reflectivity, RhoHV, deformation, divergence or vertical velocity. Time scale and height are selectable.

You can select the display style, parameters, plot limits, and so on in the **Output Options** menu of the **Quick Look** window or **Product Output Menu**.

Thresholding in the displays is based on the goodness of fit of the velocity. In **Output Options**, you can set the threshold level between 0 (display all data) and 1 (threshold all data).

More Information

‣ VVP Output Options: Line Graphs (page 260)
‣ VVP Output Options: Time-Height (page 258)

3.25 **WARN**: Warning/Centroid Product

The **WARN** product uses other IRIS products to detect significant weather.

Example: Detecting Hail

The occurrence of 45 dBZ at 1.5 km above the freezing level is a good indicator of hail in many mid-latitude locations. Assuming that the freezing level is at 4 km, and you run an echo **TOPS** product for the 45 dBZ contour, your pre-configured warning could check if:

- The **TOPS** product shows 45 dBZ tops at heights greater than 5.5 km. If yes, there is a high probability of hail.
- To avoid issuing an alarm based on a single pixel, a "threshold region" parameter checks if the region of hail signature at least 10 km^2 .
- The **VIL** for the same region (1 ... 10 km) is greater than 5 mm (or a value determined from the local climatology of hail).

The **WARN** product automates this procedure in real time by searching the products for significant weather, and alerts the operator when an event is detected. The following figure shows how the **WARN** product works.

Figure 37   HAIL Warning/Centroid

1   Threshold the input product (45 dBZ **TOPS** in the example) so that only points greater than the threshold are considered (for example, >>5.5 km (3.4 mi)).
The result is a 2-D binary array.

2   Smooth and connect the regions that are almost touching, and eliminate any isolated bins.

3   Contiguous regions are identified by a region finder procedure.
The centroid location and area of each region is computed.
Regions below the threshold size are discarded.

4   Determine whether any part of any region is in a protected area.
If so, a warning message and audio beep are broadcast to all IRIS terminals.

5   Produce the **WARN**, a situation display, also called an overlay display because it can be overlaid on any X-Y product display as selected in the **Product Output** menu.
For example, a **HAIL** warning display can be overlaid on a display of **PPI** reflectivity to show the relationship between the warning and the current weather.

Warning Message

You can define the content of the message. For example:

```
2 HAIL Warnings at 11:30:00 in: AREA_A AREA_B
```

In this case, `HAIL` is the user-selected warning text and `AREA_A` is the user-selected name of the protected area.

The names and locations of protected areas are defined in the IRIS **Setup** utility.

The messages are added to the IRIS **Message Summary** menu.

Situation Display

The contents of the situation displays are as follows:

- Outlines of active protected areas.
- Warning text at each centroid location which exceeds threshold size (such as, `HAIL` or `MBW`).
- A warning message.

The following figure shows an example of a microburst warning (`MBW`) based on the optional radial shear product. A warning display can also be transmitted with a selectable geopolitical overlay.

Figure 38   Example of Warning Situation Display

The outlines of the protected areas are shown for the areas included in the **WARN** product. In the example, the runway outlines are from a separate overlay map which is merged with the product when it is displayed.

The warning text (`MBW` in the example) is displayed at the centroid location of the weather phenomenon that has been detected. If a weather phenomenon has been detected in a protected area, the warning message is displayed in large characters across the bottom of the screen. The time associated with each warning is also displayed. If more than one hit is detected in a protected area, then the number of hits is displayed. A legend on the right of the display summarizes the characteristics of the product. This includes the product ID, date, data time and a summary of the product configuration parameters.

Warning Criteria

The automatic warning feature can provide alerts for a wide variety of weather phenomena, such as the approach of a severe storm, turbulence, lightning hazard, or flood potential.

Up to 3 criteria can be used. The thresholding and smoothing is performed separately for each, then the results are ANDed together so that centroid definition is performed on the combined field. For example, the additional criterion of the 1 ... 10 km (0.6 ... 6.2 mi) `VIL` >>5 mm (0.2 in) could be added to reduce a `HAIL` warning false alarm rate.

The following table shows some examples of warning criteria. Each criterion, surrounded by square brackets above, is one task. The results of multiple tasks are ANDed together.

Table 13   Warning Criteria Examples

Criteria	Example
Wind shear detection	```[Shear >10 m/s/km at 0.5° EL]
AND
[... at 0.7° EL]``` over an area of 3 km^2 (1.2 mi^2)
Storm turbulence detection	```[Spectrum Width >6 m/s (>19 ft 8 in / s)]
AND
[Reflectivity >20 dBZ]``` over an area of 10 km^2 (3.9 mi^2)
Hail detection	```[45 dBZ TOPS >1.5 km (>0.9 mi) above freezing level]``` over an area of 10 km^2 (3.9 mi^2)
Precipitation surveillance detection	```[1.5 to 14 km (0.9 to 8.7 mi) VIL >1 mm (>0``` $\frac{3}{64}$ ``` in)]``` over an area of 10 km^2 (3.9 mi^2)
Severe storm detection or lightning hazard	```[1.5 to 15 km (0.9 to 9.3 mi) VIL >10 mm (>0``` $\frac{25}{64}$ ``` in)]
AND
[10 dBZ TOPS >8 km (>5.0 mi)]``` over an area of 10 km^2 (3.9 mi^2)
Flash flood warning	```[Hourly Rainfall or N-Hour Rainfall >5 mm (>0``` $\frac{13}{64}$ ``` in)]``` over an area of 25 km^2 (9.7 mi^2)

> To function effectively, a **WARN** product must be based on the local climatology and experience. Vaisala can work with you to develop such a climatology or to better understand the capabilities and limitations of the **WARN** product.
>
> Vaisala makes no warranty, either express or implied, that the **WARN** product can detect all hazardous weather situations. In no event can Vaisala be held liable for damages of any kind for failure of the **WARN** product to issue a warning, or for false alarms that may be issued by the **WARN** product.

More Information

‣ WARN Algorithm (page 142)

3.25.1 Configuring WARN Products

Use the WARN product configuration menu to specify the message, the area of the threshold region, up to 3 products to use as criterion for the warning, and up to 16 protected areas.

1. In the main menu bar, select **Menus > Product Configuration**.

2. In the main menu bar, select **Type > WARN**.
   The **WARN Product Configuration** menu opens.

3. In **Warning Symbol**, specify the text used in the warning messages and the text
    displayed at the centroid position on the situation display.
    For example, the text may be, `HAIL` or `MBW`, `S++`, or `TRW+`.

4. Enter the minimum size of a thresholded region in the **Area in Sq Km** field.
    Areas that do not meet or exceed this size are discarded. Enter the desired value in
    square km.
    For example, for a 3 km by 3 km (1.9 × 1.9 mi) size area, enter **9**.

5. Select the button next to **Product Type** and **Product Name**, select up to 3 products to be examined by the WARN product.

> * Products must have the same maximum product range per the respective **Product Configuration** menus.
> * Products must have the same resolution per the respective **Product Configuration** menus.
> * Products must be of a supported data type: dBT, dBZ, dBZc, Height, Kdp, LDRH, LDRV, R, Rain, RhoH, RhoV, RhoHV, Shear, SQI, Time, VIL, Width, and ZDR.

Errors are reported at run time in the **Radar Status** menu.

a. Select the product type.
    The **Product Name** information is filled in automatically. You can also edit the name.

b. Select the product name.

> The list of product names shows products currently in your system. If the product you want does not show, run your system until it does show.
> Alternatively, chose a different product of the desired type and override the product name.

c. For each product, use **Time** values to use products from different tasks or different runs of the same task.
    Use the plus and minus buttons to increase or decrease the hours, minutes, or seconds or type a time value in the window.
    This field applies only when there is more than one criterion. For example, if the second criterion has a time of 00:10:00, when the first product becomes available, the scheduling algorithm searches back in time as far as 10 minutes to find a version of the second product.
    You must know your task schedule. If you use products based on different tasks, you must set the **Time** field to some nonzero number or the product does not run. In general, if all of your product criteria are based on the same task, set all the times to 00:00:00 so that only data from the same run are used.

d. For each product, enter the warning threshold in **Threshold**.
    The WARN product considers only those values that meet or exceed the threshold.
    The units of measure depend on the selected product. For example
    * TOPS thresholds are specified in km
    * VIL thresholds are specified in mm.

> Check the appropriate **Product Configuration** menu if you are uncertain about the units of measure.

The WARN product operates on the color bands in the product image and the threshold is set accordingly.
For example, to set up a reflectivity warning based on a CAPPI product.
- Display an example of the product.
- Choose your threshold to match the number displayed with the color band that you want to use as your threshold.

If 45 dBZ is the number displayed with the red color, and you set your threshold to 45 dBZ, then all of the regions displayed as red (or "greater") are included in the warning algorithm. This is convenient for verifying the algorithm. If you display the reflectivity product along with the situation display (warning overlay), the centroid labels are positioned precisely in the red-colored regions.
For the VVP product, the threshold refers to the divergence in units of m/s per km (10^{-3} s^{-1}). When the WARN product is run for the VVP, a warning is generated if the divergence exceeds this value at any height in the VVP. Strong low-level divergence over the radar could be an indicator of a microburst. To properly set for microburst alerts, you must know the local characteristics of the phenomena.

6. Under **Protected Areas for Warning Alert**, select **Protected Areas** to see a list of areas. The protected areas are configured in the **setup** utility. They are constrained to be rectangles with arbitrary size and orientation angle. There is a limit of 16 areas. (If you make a change to **Setup**, you must restart IRIS for the change to take effect.)

> ℹ It is a good idea to make a large area (for example, 500 by 500 km (310.7 × 310.7 mi)) named ALL. When ALL is selected, the entire radar area out to 250 km (155.3 mi) range is a protected area.

7. Select **TDWR Style** to use the TDWR warning message format.
In TDWR format, only the strongest centroid in the highest priority area is reported along with its strength. For example: MBA 3MF 30K–, in spoken language: Microburst Alert, Three Mile Final, 30 Knot loss.
The older IRIS format reports all centroids in all protected areas, for example: 3 MBA warnings at 11:30 in: 3MF, 2MF. These messages are displayed at the bottom of the display screen and optionally spoken.

8. Select **Say/Beep Warning** to have the product generator to audibly signal the warning message in addition to displaying it as a message.
You must also select **Enable Speech** from the IRIS **Setup** menu.

9. Select **Make Diagnostic** for the product generator to create additional diagnostic output.
The diagnostic is a USER product type, with a USER data type, and the same name as the WARN product, except for a MSK.
You can display this product in the IRIS Radar **Quick-Look Window**.
This product contains the output of the initial thresholded input data. The LSB is the first input, bit 1 is the second, and bit 2 is the third. For a WARN of only one input, expect a binary file. You can use this to see the regions being processed.

3.26 **WIND**: Wind Speed and Direction

wind   WIND Product Configuration: DEFAULT

File Menus Type Commands   **Help**

TASK SUMMARY

TASK Name	PPI_VOL_A	**DSP Data**	Z V W
Scan Mode	PPI Full	**Max Range**	256.0
Angle List	AZ: Full Circle	EL: 2 angles from 0.5 to 1.2	

Min–Max Rng	5.0 200.0	**Defaults**
Min–Max Height	0.5 1.5	

Range Spacing	19.5 km	**Sector Length**	19.5 km
Azimuth Spacing	45.0 deg	**Sector Width**	19.5 km

Figure 39   **WIND** Example Configuration

WIND displays wind speed and direction with either wind bards or wind strings.

WIND computes a 2-D array of horizontal wind vectors (the horizontal wind field) using the radial velocity information and the assumption that the wind is uniform over a limited sector (for example, 10 km (6.2 mi) by 60°). By computing the mean wind in a number of sectors, an approximation to the 2-D field of horizontal winds can be made. Because the algorithm assumes uniform winds over a sector, it is also called the "sector uniform wind algorithm." The algorithm can show gradual changes in the wind vectors over the radar coverage area. It cannot show sharp gradients such as fronts or microbursts.

The output can be displayed as wind barbs or wind strings as an overlay product. The mean wind can be subtracted at output so that the perturbation wind is displayed. These selections are made when the product is displayed in the **Quick Look Window** or the **Product Output Menu** using the **Output Options**.

> If there are insufficient radial wind points in a sector, the algorithm still computes a mean wind. However, the estimate is unreliable.
> For each sector, the total number of valid velocity range bins found is stored, as well as the total number expected for full coverage.
> When displaying the results, the output options allow thresholding by the percentage of coverage in a sector.
> Use this threshold to reduce noisy values. A value of 40 % areal coverage is a good starting point.

1. Select **Type > WIND**.

2. In **Min-Max Rng** specify the radius of the product, starting at a point close to the radar and ending at some distance away from the radar.
   Wind vectors are computed only within this range.

3. In **Min-Max Height** define the layer from which the radial velocity information is taken to produce the wind vectors.
   Because vertical shear can be substantial, Vaisala recommends that this layer be ~1 km (~0.6 mi).

4. In **Range** and **Azimuth Spacing**, define the resolution for computing the wind vectors in polar coordinates.
   For example, if the **Range Spacing** is set to 10 km (6.2 mi) and the **Azimuth Spacing** is set to 45°, a wind vector is computed over the **Min-Max Range** every 45° in azimuth and 10 km (6.2 mi) in range.

5. Set the **Sector Length** and **Width** values.
   Each wind vector is computed at the center of a sector which is defined by the **Sector Length** (or range) in km and the sector width (or azimuth) in degrees.
   You can set the **Range** and **Azimuth Spacing** so the sectors overlap.
   Vaisala recommends that the sector width be ~60°.

6. In **Defaults**, set all fields in this menu to their default values.

More Information

- Display Parameters (page 29)
- Product Parameters (page 25)
- Associating Products with Tasks (page 19)
- WIND and FCAST Output Options (page 264)

3.27 **XSECT**: Cross Section

Figure 40   **XSECT** Configuration Example

XSECT is a volume scan product that shows the height cross-section along a user-definable line.

The **XSECT** product is similar to an **RHI** in that it shows a range-height representation or vertical "slice" of a radar parameter.

For **XSECT**, the cross section is constructed from a **PPI** volume scan. This means that if you are running a standard volume scan at regular intervals, an **XSECT** can be constructed without scheduling an **RHI** task.

Unlike **RHI** products, which are typically composed of about 100 elevation angles, the resolution in an **XSECT** is limited by the number of elevation angles in the volume scan, typically 10 ... 20. However, the 3-D interpolation used by the **XSECT** algorithm yields an excellent representation of the height distribution of radar parameters. The algorithm also corrects for earth curvature effects so that the height scale corresponds to the actual height above your product reference height (usually sea level).

You can use the **XSECT** product in the following ways:

- For regions of interest, such as a watershed area or air terminal, you may configure an **XSECT** as a standard product to slice through the area.

- On an ad hoc basis to cut through a line of storms or an area of suspected severe weather.
  For ad hoc cross sections, use the **XSECT** tool on the **Quick Look Window** to configure the cross section line with a mouse.
  If you do not use the mouse to define a cross section line, display a **PPI** or **CAPPI** product with range rings and range numbers to help you configure the cross section.

The following figure shows the **XSECT** corresponding geometry for the previous configuration example.
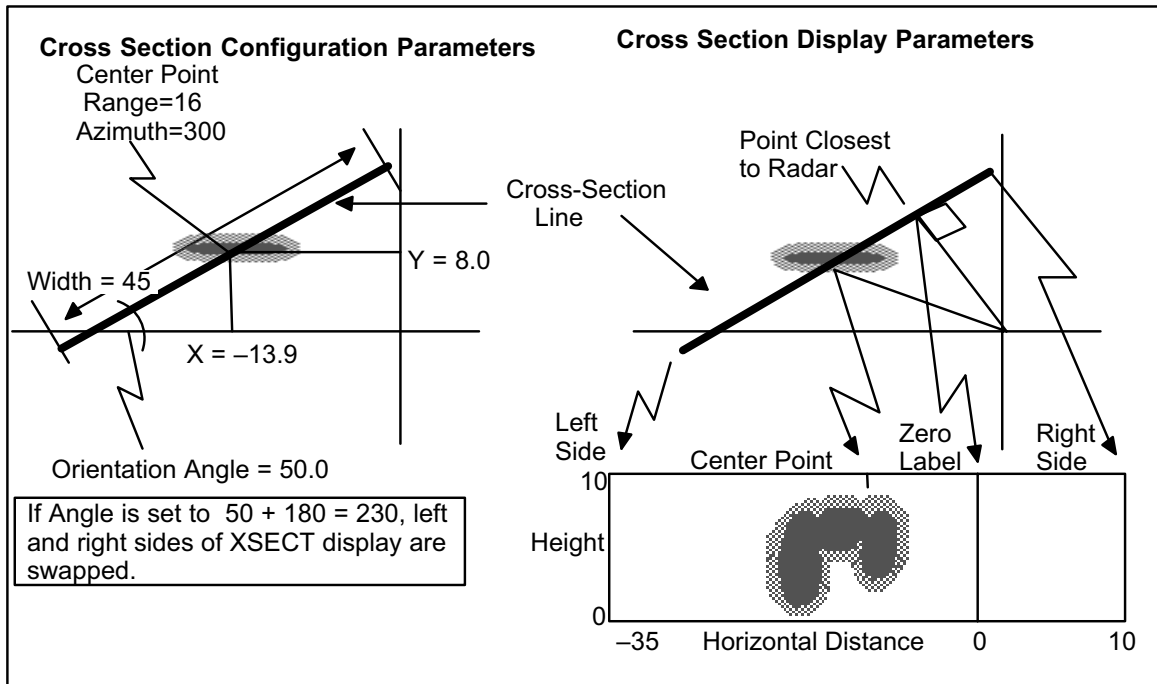
Figure 41   Cross Section Geometry

▶ 1. Select **Type > XSECT**.

2. In **Data:Display**, choose from: `Z dBZ`, `Z Rain`, `V Vel`, `W Width`, `T dBZt`, `T Rain`, `Zd ZDR`, and `Zd Rain`.

3. For **Center (x,y)**:

   a. Look at the **CAPPI** or **PPI** image that you want to "slice" and decide where you want to make your cross section.

   b. Select the center point for the cross section.
      This is where the range marks and azimuth lines are useful.
      Usually, you place the center point on the radar feature that you want to "slice," as shown in the example.

   c. You enter the coordinates in X-distance (east) and Y-distance (north) in km.
      Enter South and West distances as negative numbers.

4. Select the maximum height and the horizontal width of the cross section in km and enter them in **Height Width**.
   There is no need to be exact, use the range and azimuth marks that are on the display to estimate the horizontal width.
   In the example, the height is 10 km (0.6 mi) and the width is 45 km (28.0 mi).

5. In **Angle**, select the orientation angle in degrees.
   In the example, the angle is 50°. The angle that you select determines the end of the cross section slice that is displayed on the right side of the display. An angle of 50° displays as follows:

   ```
   Left Side          Right Side
   23° end            50° end
   ```

   To reverse the ends in the output display, select an angle of 50° + 180° = 230°.

6. In **XZ Smoother**, set the X and Z smoothing length scales independently.

   > ℹ️ Because of the interpolation performed in constructing the cross section product, the output displays usually require little, if any, smoothing.

   In the final output display, the zero point of the horizontal range scale corresponds to the point on the cross section closest to the radar. This is also shown schematically in the example.
   This means that if you select a line that goes through the radar (which is allowed to make a pseudo **RHI**), the 0 point is at the radar, so the horizontal distance equals the radar range.

More Information

‣ RHI: Range Height Indicator (page 68)
‣ Display Parameters (page 29)
‣ Product Parameters (page 25)
‣ Associating Products with Tasks (page 19)
‣ Should I Use RHI or XSECT? (page 71)

4. Scheduling Products

4.1 Product Scheduler Menu

The **Product Scheduler** menu tells IRIS when to generate a product for automatic or manual operation.

Because the product generation affects host computer resources, this menu can be viewed but not controlled by an IRIS observer.

Figure 42   **Product Scheduler**

Site column
> Shows the site for each scheduled product. A product runs only if data from that site comes to IRIS, either from the radar, over the network, or from a tape.
> If the **Display** field is set to **Master** (the **Master Schedule**), there is only one product entry, regardless of the number of sites scheduled.
> Right-click the mouse to show a menu on the site column listing all sites with a * symbol next to each site for which the product is scheduled.
> Select an entry to add or remove the * tag. This makes the **Master Schedule** convenient to use for adding or deleting sites from the schedule.

Product column
> Shows the name of the product.

Data column
> Shows the data parameter configured for the product.

Task column
> Shows the name of the associated task.

Next-Data-Time

Scheduling operates under the concept of data time, the time when a task starts collecting data. For an on-time task schedule, the data time corresponds to the time when a task is scheduled to start.

When you schedule a product, set the **Next-Data-Time** field for the product generator to process only associated tasks with data times later than the next data time.

Skip

The **Skip** time breaks the day into equal intervals of time starting from midnight. A product is generated, at most, once for each skip time interval. The first occurrence of data from the associated task is used for each interval.

Request

Schedules a product according to whether the schedule affects **All**, **Next** occurrences of the product. You can also stop the product's schedule.

Status

Shows the current status of each product as **Running** (product is being generated) or **Wait** (product is waiting for either the associated task to run, for its turn to run pending the completion of other products, or the product is not scheduled to run). After a product is generated, the status changes from **Running** to **Wait**.

When the **Master Schedule** is displayed, **Status** shows the status of the site which is either currently running or ran most recently.

Runs

A counter (`000` to `999`) showing how many times the product has run since it was loaded in the menu.

If the number of runs exceeds `999`, the counter restarts at `000`.

If you stop a product, the counter is not reset. The counter resets if you delete a product and reload it into the schedule. If you load a new schedule, the counters for all products are reset.

When the **Master Schedule** is displayed, the **Runs** field reflects the total number of runs for all sites that are scheduled.

4.2 Scheduling Products

▶ 1. Select **Menu > Product Scheduler**.
A list of available product types appears.

2. In **Display**, on the top of the menu, specify which sites are displayed.

Value	Description
All Sites	Show the schedule for all sites.
Master	Show the **Master Schedule**, with one entry for each scheduled product.
	If a product is scheduled to run at multiple sites, the entry shows the site that ran most recently.
Site Group	Show the products that are configured for all sites in the group.
	The **Site Group** is a subset of sites that is defined in the **Site Group** field.
Single Sites	Show the products configured for a single site (select the site name). Use this if your system supports one site.

The adjacent text field shows how many products match the display selection.

3. In **Add for**, on the top of the menu, specify the sites for which the product is added.

Value	Description
All Sites	Add the product for all sites.
Site Group	Add the product for all sites defined by the **Site Group**.
	Tag with * the members of the site group. If you have a single site system, make your site the only member of the **Site Group**.
Single Sites	Add the product only for that site. Use this if your system supports one site.

4. To edit the products configured in the product schedule, right-click the **Product** field and select one of the following:

Value	Description
Add	Add a product to the schedule.
	The product is generated on an on-going basis when the task collects data, subject to the **Skip** time.
Remove	Remove the selected product from the list for the site that is displayed.
	If the **Master Schedule** is displayed, then the product is removed for all sites.
Edit	Opens the **Product Configuration** menu for the selected product.

5. To generate data only from future tasks:

a. Right-click the **Request** field

b. Select **All** or **Next**.

Value	Description
All	All associated task data collected after the next data time are processed.
	The product is generated on an on-going basis when the task collects data, subject to the **Skip** time.
Next	Only data from the next occurrence of the task, after the next data time, are processed.
	The product is generated only once.
	The **Skip** time is ignored.

6. If you use **All** requests, you can set the **Skip** field so that the product generator does not process every occurrence of the associated task.

   By default, an **All** request causes data from every task to be processed. The **Skip** time breaks the day into equal time intervals starting from midnight.

   A product is generated, at most, once for each skip time interval. The first occurrence of data from the associated task is used for each interval.

   For example, a **Skip** time of `00:15` (15 minutes) means that the product is generated no more frequently than every 15 minutes, regardless of how often the associated task runs.

   In the following figure, the **Skip** field for an **All** request is set so that there are 2 occurrences of the task in each interval. Only the first occurrence is used to generate the product.
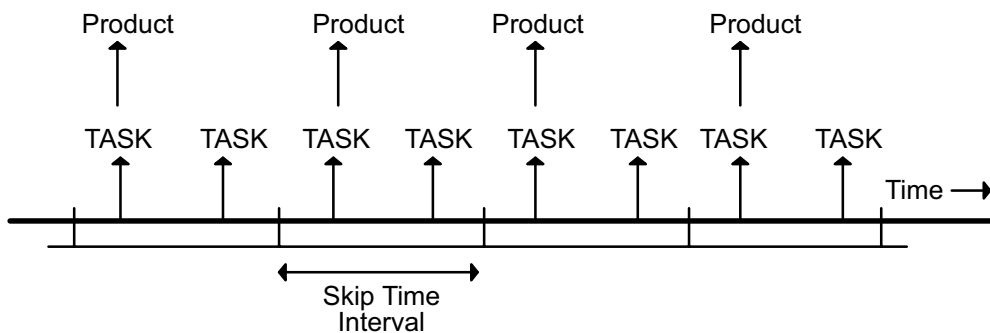
   The default value of the **Skip** field is `00:00`, indicating no tasks are skipped.

   When the **Master Schedule** is displayed, a change in the **Skip Time** is applied to all of the sites.

   > You must consider the task schedule when specifying the **Skip** time.

   During operation, the **Next-Data-Time** field changes to show the next possible time that the product can be generated. This depends on the **Skip** time and the request (**All** or **Next**) as follows:

Skip and Request Values	Description
Next Any Skip Time	After completing a **Next** request, the **Next-Data-Time** field is reset to the data time (plus 1 second) of the task that was just processed.
	This means that issuing another **Next** request processes the next occurrence of the associated task.
	The **Skip** time is ignored.
All Skip Time = 00:00	The product is generated for every occurrence of the associated task.
	The **Next-Data-Time** field is reset to the data time (plus 1 second) of the task that was just processed (identical to the **Next** request case).
	This is the same as issuing a series of **Next** requests.

Skip and Request Values	Description
All Skip Time >00:00	The **Skip** field determines the earliest time that a new product can be generated.
	After completing a product, the **Next-Data-Time** field is reset to the beginning time of the next skip time interval.

7. To generate data only from future tasks, right-click the **Request** field.

Request Field Options	Description
All	All associated task data collected after the next data time are processed.
	The product is generated on an on-going basis when the task collects data, subject to the **Skip** time.
Next	Only data from the next occurrence of the task, after the next data time, are processed.
	The product is generated only once.
	The **Skip** time is ignored.
Stop	Stops the selected product's schedule.

When the **Master Schedule** is displayed, a change in the **Request** is applied to all of the sites.

8. To generate products from old data, right-click the **Next-Data-Time** field.
This allows you to generate products from task data received from another IRIS host or retrieved from tape.
The following example shows how **Next-Data-Time** interacts with **All** and **Next** request options.
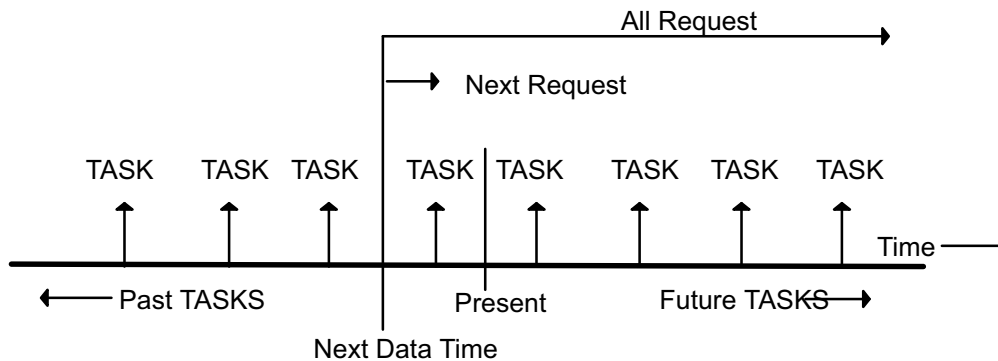
When a product is first loaded into the schedule, or an entire new schedule is loaded, the **Next-Data-Time** field defaults to the current time.

a. In the **Next-Data-Time** column, right-click the time stamp.

b. Select **Next > Product > Time buttons** to go back in time.

9. Select **File > Save** and give a name for the product scheduler.

4.3 Adding a Product to Schedule

▶ 1. Select the header or product for the type of product you want to enter.

2. Position the cursor over the **Product** field and select **Add**.
    IRIS displays a list of products of that type.

3. Select a product.
    The product is added to the schedule.

4.4 Removing a Product from Schedule

▶ 1. Select the product you want to remove.

2. Place the cursor over the **Product** field and select **Remove**.
    You can remove products only. IRIS does not remove product headers from the
    schedule.

4.5 Editing Product Configuration

▶ 1. Select the product you want to edit.

2. Position the mouse cursor over the **Product** field and choose **Edit**.
    IRIS opens the **Product Configuration** menu with the selected product loaded into it.

3. Edit the product as needed.

4. Select **File > Save**.

5. Exit the **Product Configuration** menu.
    IRIS returns you to the updated **Product Scheduler** menu.

4.6 Stopping Product Generation

▶ 1. Select the product you want to stop.

2. Position the mouse over the **Product** field and choose **Stop**.
    If the product is waiting to be generated, the **Stop** request unschedules the product
    from the **All** or **Next** states.
    If the product is being generated when you make the **Stop** request, no product output
    file is produced.

4.7 Tips for Running Radar Products

Optimize System Performance

You can configure an almost unlimited number of products in IRIS. The system manager must carefully plan the product mix and performance capabilities for your installation.

Table 14   System Performance Considerations for Generating Radar Products

Consideration	Description
CPU capabilities	You can improve performance by off-loading product output processes to a different machine. Consider generating products on a separate IRIS Analysis system.
Product mix	Avoid wasting computer resources by making products that no one will look at. For example, people do not usually look at 12 **CAPPI** products. A mix of **CAPPI**, **VIL**, and **TOPS** may be more appropriate.
Output devices	Consider the trade-off between output devices and product mix. If you are driving many remote nodes, you can generate fewer products.
High resolution displays	Do not create if not required. Medium and low resolution products run more quickly. Prefer low resolution products for serial line remote displays because the update rates are faster.
High resolution sampling	Do not use in your task configurations if not required. If you want to use high resolution sampling (in azimuth and range), do not sample at ranges greater than you need for your application.
Custom products	If users want to request many custom products, consider purchasing a separate workstation for running display software. This means you can to obtain **RAW** data from the IRIS host and process products on the separate workstation.

Immediate Products

Most IRIS products are volume scan products that require a series of **PPI** products.

Immediate products can run before a volume scan is complete because they are based on a single sweep of data. Use immediate products for applications where rapid feedback is required. For example, **PPI**, **RHI**, or **SHEAR**.

Making Products from Retrieved Data and Keep Flags

When **RAW** product files are restored from a tape, they are automatically reingested to reconstruct the ingest files required for product generation.

Because the data from tape are old, they are first on the list to be deleted by the system Watchdog process. To avoid this, do the following in the **Ingest Summary** menu:

- Before retrieving files from a tape, delete some unneeded ingest files.
- After the files are restored, tag the reingested files with `Keep` flags so they are not deleted when real time files come in.

For information on space reserved for ingest files, see **Setup** in *IRIS and RDA Utilities Guide.*

5. Configuring and Scheduling Optional IRIS Products

5.1 CATCH: Subcatchments Precipitation Accumulation

CATCH calculates the precipitation accumulation in subcatchment areas such as watershed areas.

CATCH is used for hydrometeorological applications such as estimating the total rainfall in a river basin for the purpose of flood forecasting. Both the time of the integration and the subcatchment areas can be selected.

CATCH can also issue warnings if the precipitation in a subcatchment region exceeds a threshold value.

CATCH is a product of a product for which the hourly **RAIN1** product serves as the input. You can sum any number of hours of individual **RAIN1** products. The only limitation is the number of **RAIN1** products stored on disk. As part of the IRIS Hydromet option, the **RAIN1** products can be calibrated by means of raingage input, thus **CATCH** products also get the benefit of this calibration.

The subcatchment areas are defined similarly to overlays, that is, subcatchment regions are defined by using LAT/LON vector points and stored in a file in the *IRIS_OVERLAY* directory.

A "catch" file typically contains many subcatchment areas. There can be more than one catch file so that users can have different "mixes" of subcatchments.

The output of the **CATCH** product is a file that contains the precipitation depth for each subcatchment for each hour and the total over the selected integration time. This file can be displayed by itself or overlaid on another product display. An interactive bar graph display shows the hourly precipitation amounts in each subcatchment.

> You must purchase a license for the Hydromet option to be able to use the **CATCH** product.
> If the **CATCH** product is not listed in the **Product Scheduler** or in the list of product types in the **Product Configuration** menu, then you do not have a license for the **CATCH** product.

More Information

‣ Hydromet Raingage Correction Overview (page 377)

5.1.1 Configuring CATCH Products
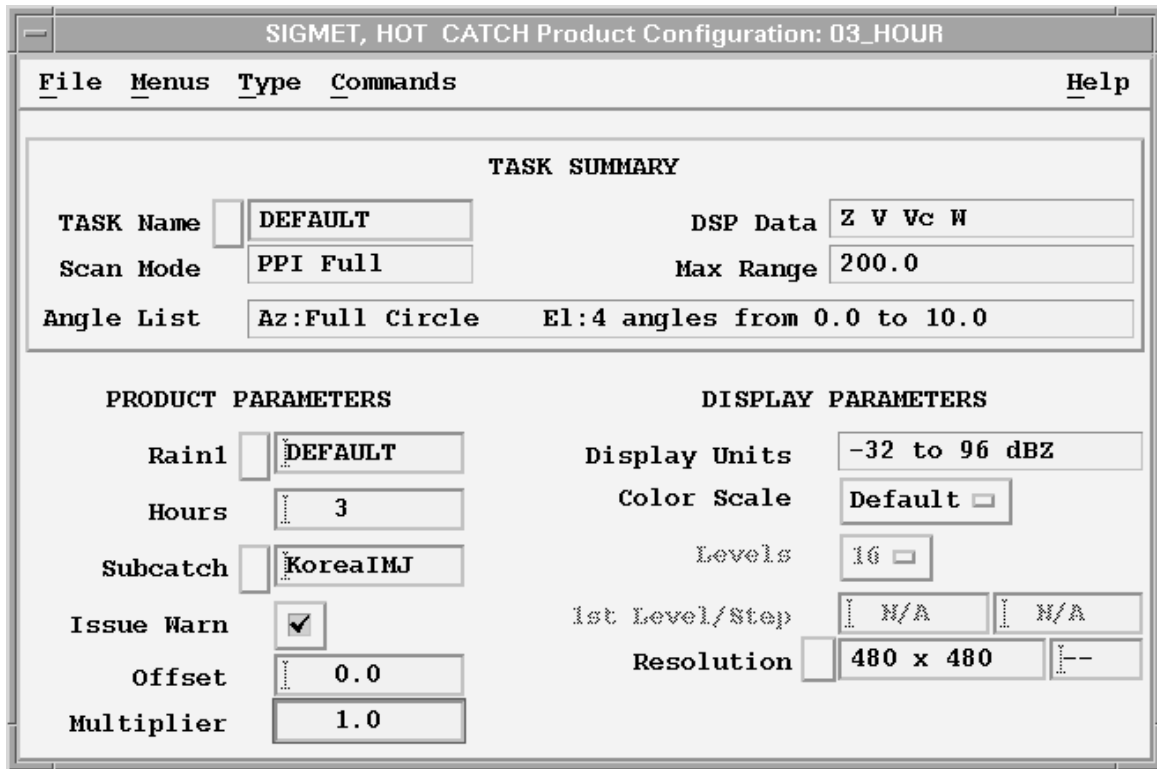
Figure 43   **CATCH** Example Configuration

▶   1. Select **Type > CATCH**.

2. Configure the **CATCH** product:
    Rain1
    > Specify the product name of the input hourly **RAIN1** product.
    >
    > Select **Rain1** and choose from a list of products, or type the name in the field.
    >
    > When you enter a product name, the associated task information is displayed in the **TASK Summary**.

    Hours
    > Specify the number of hours to integrate. When the product runs, it integrates data for the last N hours.
    >
    > If a **RAIN1** product is missing for one of the hours in the interval, the algorithm assumes that no rain fell during that hour.
    >
    > The product output shows how many hours were integrated.

    Color Scale, **Levels**, and **1st Level/Step**
    > Specify the color levels can be specified.
    >
    > The resolution is fixed by the resolution of the **RAIN1** configuration.

    Subcatch
    > Select the subcatchment file name. The subcatchment file contains the subcatchment definitions.

    Issue Warn
    > The **CATCH** product can generate warnings if the amount of precipitation for a subcatchment exceeds a threshold value.
    >
    > The threshold value for each subcatchment area is configured in the subcatchment file.

    Offset and **Multiplier**
    > Use these values to modify the thresholdsfor each subcatchment. This might be done to account for different integration times.
    >
    > For no adjustment of the thresholds in the subcatchment file set `Offset=0.0` and `Multiplier=1.0`.
    >
    > Vaisala recommends that you configure the thresholds in the subcatchment file to be characteristic of accumulation over a standard time such as one hour. You can then use The offset and multiplier can then use **Offset** and **Multiplier** to adjust all of the thresholds.
    >
    > For example, assume that an indicator of flood hazard in a subcatchment experience shows that the one hour threshold is 10 mm (0.39 in) and the 2 hour threshold is 15 mm (0.59 in) for flood hazard. You can configure the 10 mm (0.39 in) value for the subcatchment and then set the multiplier to 1.5 and the offset to 0.0. This allows some degree of adjustment of the subcatchment thresholds for different time periods.
    >
    > Another technique is to configure different subcatchment files with the thresholds set for different integration times. In this case the name of the subcatchment file should reflect the integration time such as *CATCH_06* to indicate that the thresholds are set for 6 hours.

3. In **Save As**, name your **CATCH** product.
   Vaisala recommends naming your **CATCH** product by the number of hours of integration and a general descriptor of the area, perhaps even the name of the subcatchment file.
   For example, if the subcatchment file for a large river watershed is called *RED_RIV* and the integration time is 2 hours, save the product under the name *RED_RIV_02*.

More Information

‣ Subcatchment Definition (page 130)

5.1.2 Subcatchment Definition

Subcatchment areas are defined in a file stored in the *IRIS_OVERLAY* directory (typically */usr/sigmet/config/overlay*).

These are given the suffix *.cat* and are referred to as *catch* files.

When IRIS starts, the start-up log shows the processing of the catch files.

In general the files contain the following information for each subcatchment:

- Subcatchment index number (an integer 1, 2, 3, ...) and name.
  The name is used for display.
- The threshold (mm) of accumulation depth for the warning.
- The subcatchment points in LAT/LON coordinates.
  These points must define a closed region and be in sequence about the perimeter of the region. That is, the first point and the last point must be the same.
  Subcatchments that share the same boundary must be entered as complete closed regions, that is, the shared boundary points are the same for adjacent subcatchments.
- The location of an internal point for the labeling.
  This point must be within the bounded region of the catchment.

Each catch file usually has many subcatchment regions. There can be multiple subcatchment files for different applications such as:

- For different radars in a network.
- To specify different warning thresholds to account for different integration times.
- To cover different major watershed areas.

Overlapping subcatchments in a file are permitted. For example, a large watershed region could be input as a single large catchment and the smaller tributary watersheds within it be specified as well. IRIS allows a maximum of 20 different catch files to be created.

For more information, see the **Overlay** utility in *IRIS and RDA Utilities Guide*.

More Information

‣ Configuring CATCH Products (page 128)

5.1.3 Scheduling CATCH Products

The **Data Time** that you specify when you schedule the product are quantized to the time value in the product configuration for the nearest interval that ends the integration.

Use the **Skip Time** field if you want only the accumulations for specific time intervals. You can set the skip time to any multiple of the accumulation period. For example, if you want a 3-hour integration only for the periods ending at 03:00, 06:00, 09:00, ..., set the **Skip Time** field to 03:00.

> If the **RAIN1** input is corrected by means of raingage input data, then there is a delay in the **CATCH** product since the input **RAIN1** product is delayed as it waits for the arrival of the raingage data.

Scheduling **CATCH** products is the same as scheduling **RAINN** products.

More Information

‣ RAIN-N: N-Hour Rain Accumulation (page 63)

5.1.4 CATCH Algorithm

The **CATCH** product algorithm takes the **RAIN1** products for the specified integration time (hours) and performs the following steps:

1. For each subcatchment, all pixels in the **RAIN1** product that are within or touching the subcatchment perimeter are added together and then divided by the total number of pixels.
   Thresholded pixels are treated as 0 rainfall.
   The result is the average hourly rainfall for the subcatchment.
2. This is repeated for each hour and each subcatchment.
3. The sum over all hours in the integration time is taken for each subcatchment to get a total average accumulation for the time period.
   If the **Issue Warn** button is enabled, this sum is checked against the catch file warning threshold (adjusted by the **Offset** and **Multiplier**) and a single IRIS warning is issued for all subcatchment areas that exceed the threshold.

For each subcatchment, the final output file contains:

• The LAT/LON of the label point, the point name and index number.
• The hourly average accumulations over the selected integration time.
  These are used for the bar graph display.
• The total average accumulation for the integration time.
• The fraction of pixels in the subcatchment that had non-zero precipitation.

The output file does not contain the subcatchment boundary definitions. This has implications for display when the **CATCH** product is transferred to another machine. See 5.1.5 Displaying CATCH (page 132).

5.1.5 Displaying CATCH
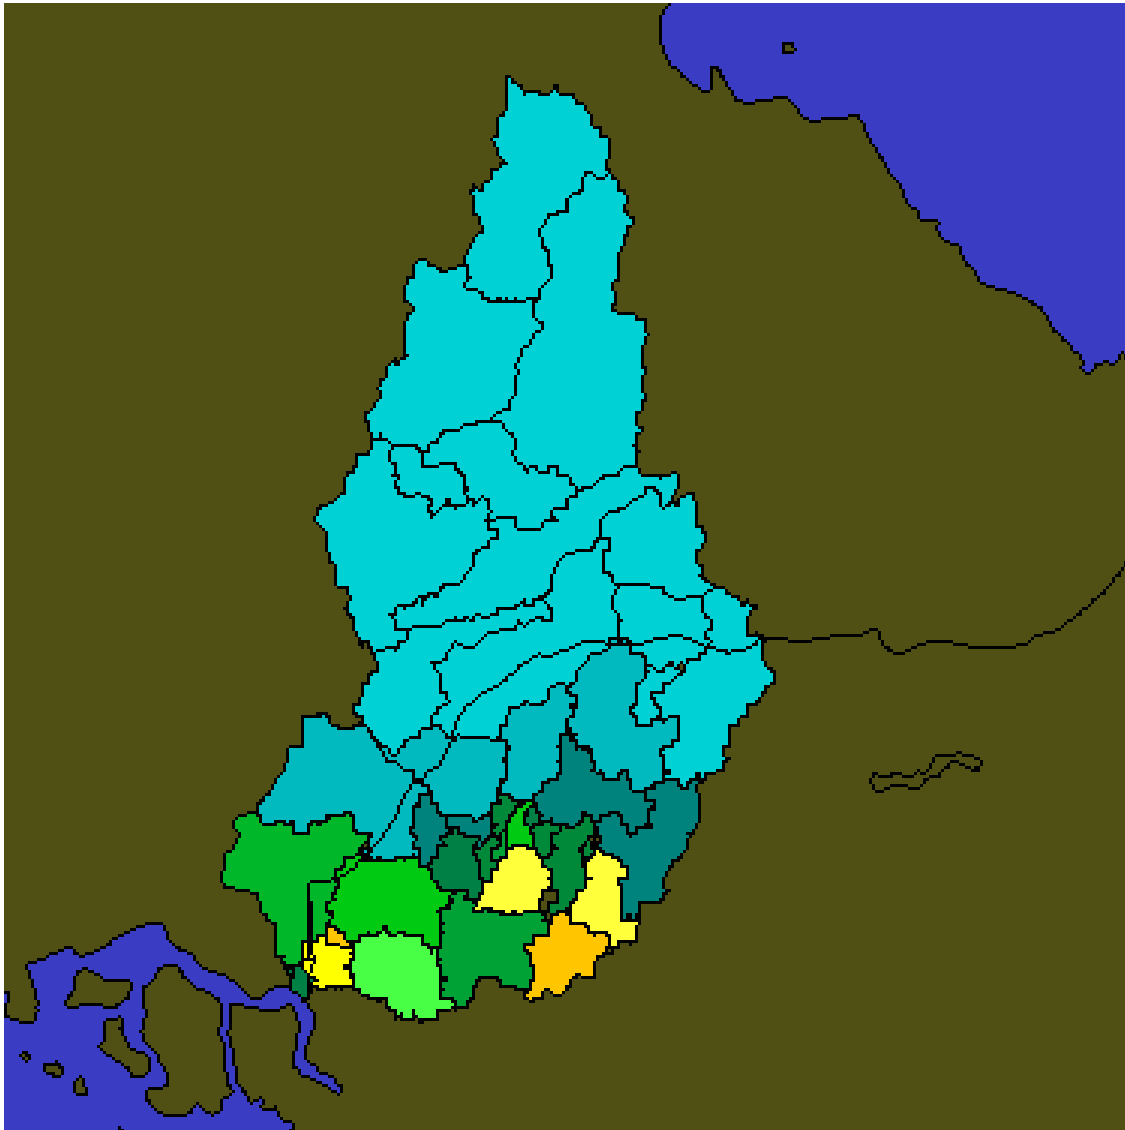
Figure 44 **CATCH** Product Display

You can display **CATCH** products in the **Quick Look Window**, as well as in other display outputs.

The **Quick Look Window** display for the **CATCH** product shows either:

- Color-coded subcatchment map regions where the color code shows the precipitation depth (see example above). The color scale is the rainfall depth scale in mm or inches.

- Colored icons where the icon (a raingage bucket shape) is drawn in the color code for the precipitation depth. This display style is used when the **CATCH** product is overlaid on another product.

1. If needed, configure the color scales and the units in the **Color Setup** utility.
See *IRIS and RDA Utilities Guide* .

2. In the **Quick Look Window**, display a **CATCH** product.

3. To choose a display style, select **Quick Look Window > Options**

4. Right-click the mouse on a subcatchment
A bar graph appears showing the hourly precipitation accumulation for the integration time.

5. To select the bar graph vertical scale, use the up and down indicators next to the top of the scale.

5.2 **COMP**: Composite
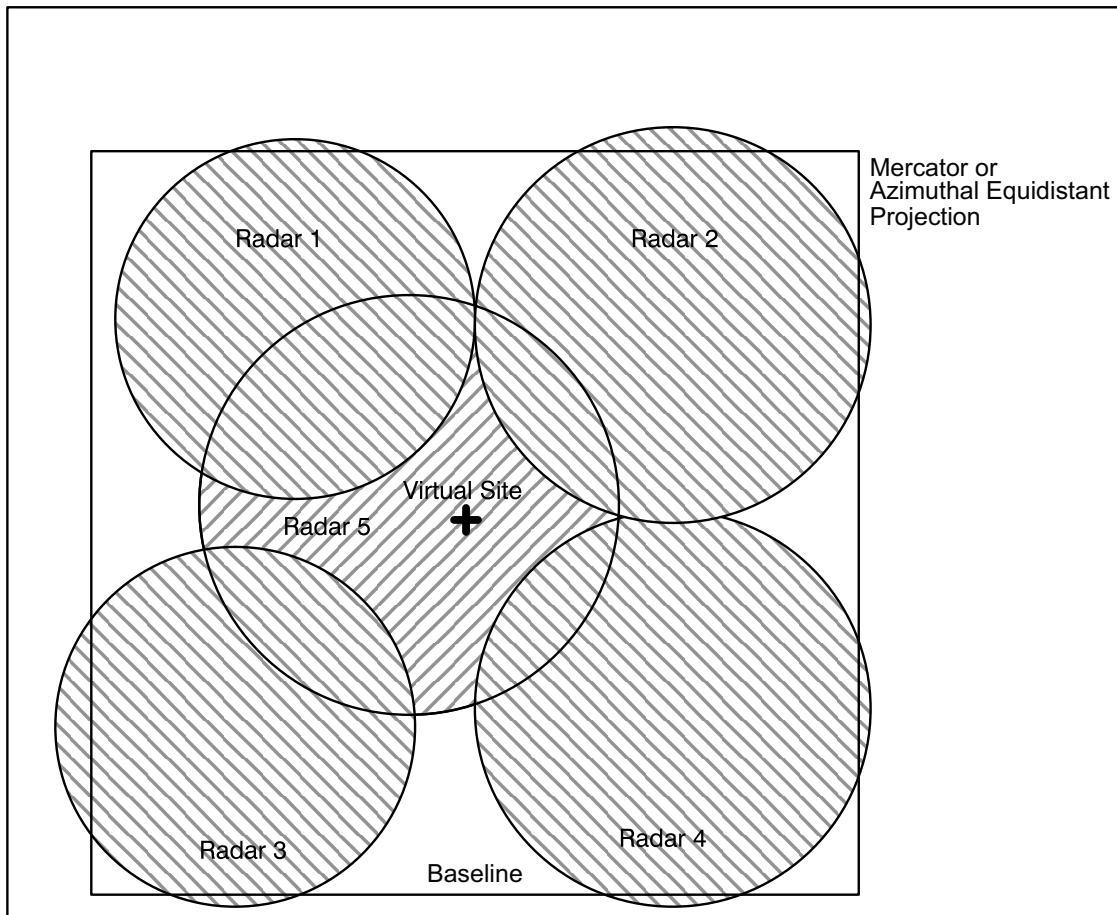
Figure 45   Example of Composite Projection Geometry

COMP produces a single composite image of radar data from many radar sites.

COMP allows you to combine data from many different radars.

The advantages of compositing are:

- Expand the area of coverage to give forecasters the big picture.
- Fill-in blind spots caused by mountains or required sector blanking.
- Fill-in blind spots caused by scan strategy limitations (for example, not scanning to high elevation angles).
- Simplify product management so users do not need to check multiple single-radar images.

Types of Products for Compositing

Unlike other products, when you schedule a **COMP** product, IRIS does not make a product of type **COMP**. Instead, it makes a composite of the input product and saves it as the input product type.

For example, if a **PPI** is the input to the **COMP** product, then the output is a **PPI** product.

The types of products that can be combined by **COMP** are:

- **BASE**
- **CAPPI**, including 3D **CAPPI** (dBZ and R)
- **HMAX**
- **MAX** (side panels are removed)
- **PPI** (dBZ and R)
- **RAIN1**
- **SHEAR**
- **SRI**
- **THICK**
- **TOPS**
- **USER**
- **VIL**
- **WARN**

Use of Composites by Other Products

COMP produces real IRIS products instead of just "pasting" the images together to make a combined picture.

This means that other IRIS products, which use products as input, can be run on the composites. For example, a composite **CAPPI** can be used as input to a **RAIN1** accumulation product. A composite 3D **CAPPI** can be used as input into the optional IRIS 3D package. Products such as **WARN**, **TRACK**, and **FCAST** can run on IRIS products that are composites.

Automatic Remapping of Input Products

Compositing of products can be complicated because of issues of different radar resolutions, maximum range, angle spacing, number of pixels, map projections, and similar.

IRIS makes it easy by automatically remapping the input products to the final output projection and resolution. This means that you can input products that have different pixel resolution, center position and maximum range, even in different projections (Mercator versus azimuthal equidistant), and IRIS does the remapping to make the final product.

This means that when you make composites, you can use normal single-radar projections. There is no need to make a special set of products to serve as input to the composites. This simplifies system configuration management.

Sources for Input Products

From the the host computer, select where input products are made and who makes them (IRIS against non-IRIS). The products used by **COMP** can be:

- Generated locally on the same workstation where the composites are made.
- Sent over the network from another workstation (IRIS or non-IRIS using UPI).
- Any combination of the above.

Input of Products from Non-IRIS Systems

The ability to send products over the network to a central compositing computer, plus the ability to automatically remap data, makes it easy to combine data from non-IRIS systems.

The remapping feature is important since you may not have control over the range or resolution of the non-IRIS data. Products from non-IRIS systems can be reformatted to match the standard IRIS formats and then inserted into IRIS using the Input pipes mechanism.

Vaisala supplies source code examples to customers who want to reformat non-IRIS data and insert it into IRIS. These products can then be sent to the composite IRIS workstation either directly, from an external user program, or through another IRIS.

Composites of Composites (Regional, National/International Composites)

The maximum number of radars that can be used in a composite is 16. This is adequate for most applications.

For larger networks, it is recommended to first make regional composites, and then composite the composites to make national or international composites.

5.2.1 Composite Algorithm

Basic Rules

COMP takes products from several radar sites and combines them into a single product using the following rules:

- The input products must be of the same type (for example, **CAPPI**) and have the same name.
  This is the recommended way to manage products on your network in general.
  Note that IRIS sorts everything by site for you, so there is no need to put a site ID in the name.
  The product type and name are specified in the **COMP Product Configuration Menu**.
- The products must be within a certain time window to be combined.
  You would not want to combine data if the input products were more than a few minutes apart.
  The time window is specified in the **COMP Product Configuration Menu**.

Remapping

The products to be composited are remapped to the common projection that is specified in the **COMP Product Configuration Menu**.

Ideally, **PPI**s and **CAPPI**s should be at the same elevation angles and heights.

If they are not at the same elevation angles and heights, there are tolerances for combining **PPI** products of slightly different elevations (for example, 1°) and **CAPPI** products of different heights (1000 m (3280 ft 10 in)).

Combining

For regions where radars overlap, you can select one of the following methods for combining radar data:

- *Maximum*
  Maximum uses the maximum value to combine the data. This is the most common setting.

- *Average*
  Average uses the average of the available data. This is a poor choice if you are trying to cover blocked regions.
- *Priority*
  Priority uses data from the available radar that is highest on the site list defined in the **COMP Product Configuration Menu**.
- *Weighted*
  For weighted, the assigned value at a location is determined by a weighted average from all the radars detecting a signal at that location.
  The weighting is inversely proportional of the distance to the radar site. Data from closer radar sites have a bigger impact to the final value than data from sites farther away.

You can choose the combining algorithm in the **COMP Product Configuration Menu**.

Output

There is no **COMP** product. The output of a **COMP PPI** is a **PPI** product with the name that is assigned in the **COMP Product Configuration Menu**.

More Information

‣ Configuring COMP (page 138)

5.2.2 System Manager Preparations for Composites

To support composite creation, the system manager must set-up the following items in advance. In many cases, these are pre-configured for your system.

- Configure the radar sites in IRIS **Setup > General**. See *IRIS and RDA Utilities Guide*.
- Configure the projection in IRIS **Setup > projection**. See 3.2.2 Configuring Map Projections (page 21).

5.2.3 Configuring COMP
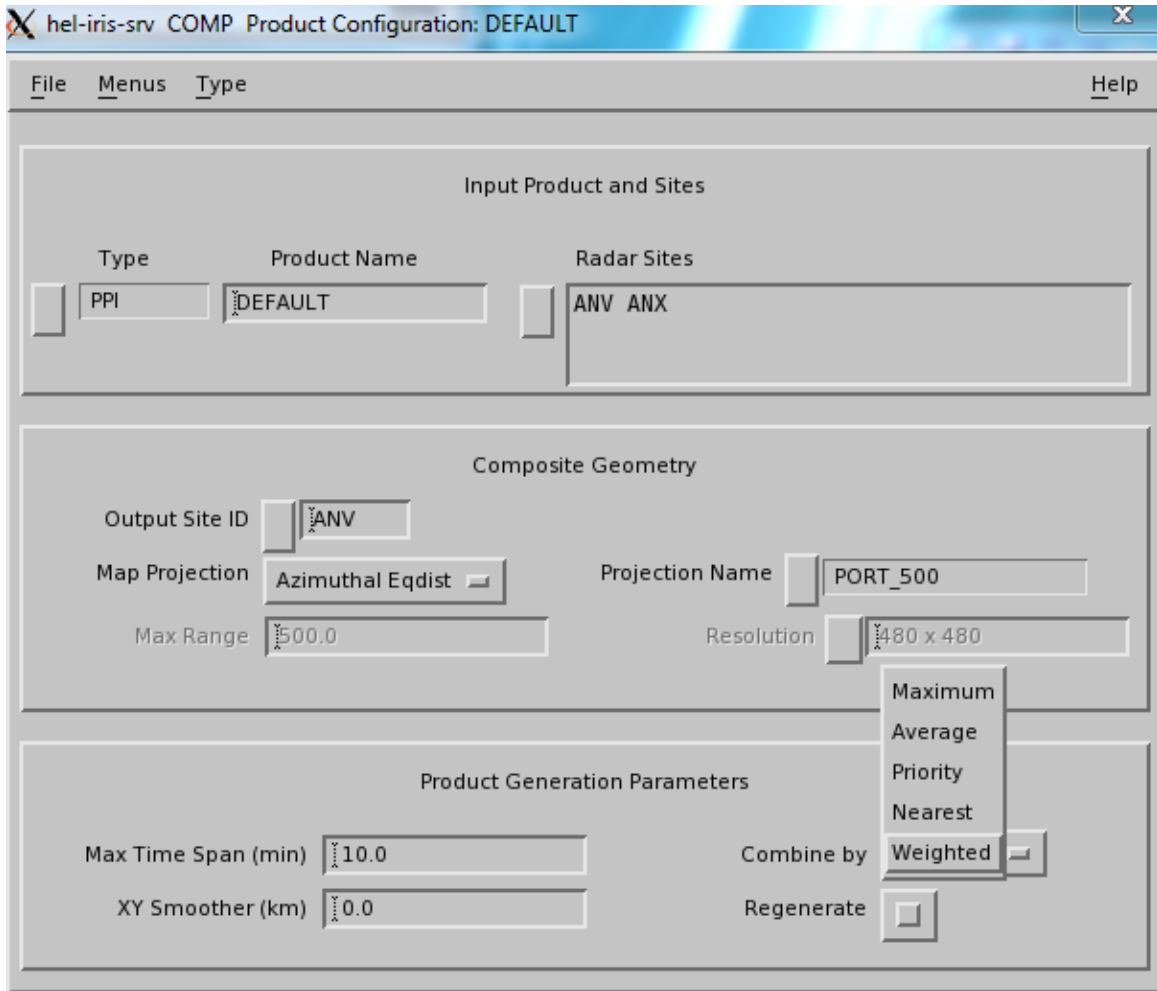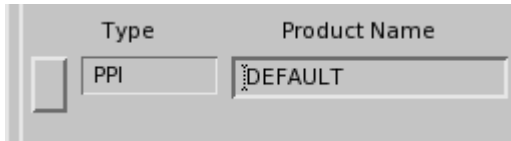
Figure 46   **COMP** Configuration Example

▶ 1. Under **Type**, select the product type and name that you want to composite.

The products from the different sites must all have this name.
IRIS only allows you to composite the following products that are licensed on the compositing computer:
- **BASE**
- **CAPPI**, including 3D **CAPPI** (dBZ and R)
- **HMAX**
- **MAX** (side panels are removed)
- **PPI** (dBZ and R)
- **RAIN1**
- **SHEAR**
- **SRI**
- **THICK**
- **TOPS**
- **USER**
- **VIL**
- **WARN**

2. Select **Radar Sites** and use the radar site selection tool to load the sites that you want to composite

> The order of the sites is important. The highest priority (first on the list) site that is in the composite (some may be missing) determines the "Data Time" that is associated with the product.
> It is also used for filling data in overlapped regions if you select the **Priority** combining rule.

3. In **Output Site ID**, specify the site used to identify the composite.

This does not have to be an actual radar site. You can use a "Virtual Site" associated with the projection or your network.
Sites are configured by your system manager in the **Setup** utility. See *IRIS and RDA Utilities Guide*.

4. In **Map Projection**, define the common projection.

Input products are remapped to this.
Typically, map projection configurations are done by your system manager.
The **Max Range** and **Resolution** fields are determined by the projection, that is, you cannot edit these.

5. In **Max Time Span (min)**, define the maximum time difference that the algorithm can allow for compositing products.

The product times for volume scan products correspond to the start of the volume scan (the Data Time).
For **PPI** it is the actual sweep time. A value of 5 minutes is probably close to the maximum that you would typically allow. If you are unsure, consider how quickly things move in a typical loop.
This number must be less the period between volume scans in your input data.

6. In **Output Smoother**, add additional smoothing if needed.

In many cases, the input products are smoothed.

> Smoothing is computationally intensive, especially for 3D **CAPPI** cubes. If your workstation is too busy, you may want to reduce the smoothing.

7. In **Combining Algorithm**, select the algorithm that is used in regions where the radars are overlapped.
    The options are **Maximum**, **Average**, **Priority**, **Nearest**, or **Weighted** (per the site list).

> **Maximum** is recommended for filling blanked or blocked regions.

More Information

- Configuring Map Projections (page 21)
- Composite Algorithm (page 136)

5.2.4 WARN Algorithm

The **WARN** product uses a composite algorithm.

Remapping

WARN products are not stored in Cartesian format, so there is no remapping required. The output projection specified in the **COMP** configuration determines the default display location and range when the composited **WARN** product is displayed.

Overlapping

The key to combining **WARN** products is determining which of the input centroids are really the same feature, and which are different. We do this with the adjustable **Centroid Distance** specified in the **Product Configuration Menu**.

For **WARN** products, this field replaces the **Output Smoother** field used for other products.

The composite is started by copying the highest priority input. Then all the other inputs are looped through in priority order. Each new centroid is compared with all the centroids in the composite being assembled. If the center of a new centroid is within the centroid distance of the ellipse of a composite centroid, then they are treated as overlapped in the combination algorithm below.

Combining

The choices for combining algorithms in the **Product Configuration Menu** (see 5.2.3 Configuring COMP (page 138)) are different for **WARN**:

- **OR Max**

  The output contains a centroid for each centroid in any of the inputs. If 2 centroids overlap, then the maximum strength centroid is kept.
- **OR Avg**

  The output contains a centroid for each centroid in any of the inputs. If 2 centroids overlap, then the centroids are averaged together.
- **AND Max**

  The output contains a centroid only if all the inputs have a centroid at that location. The maximum strength centroid is used.
- **AND Avg**

  The output contains a centroid only if all the inputs have a centroid at that location. The average value centroid is used.

More Information

‣ WARN: Warning/Centroid Product (page 104)

5.2.5 COMP Scheduling
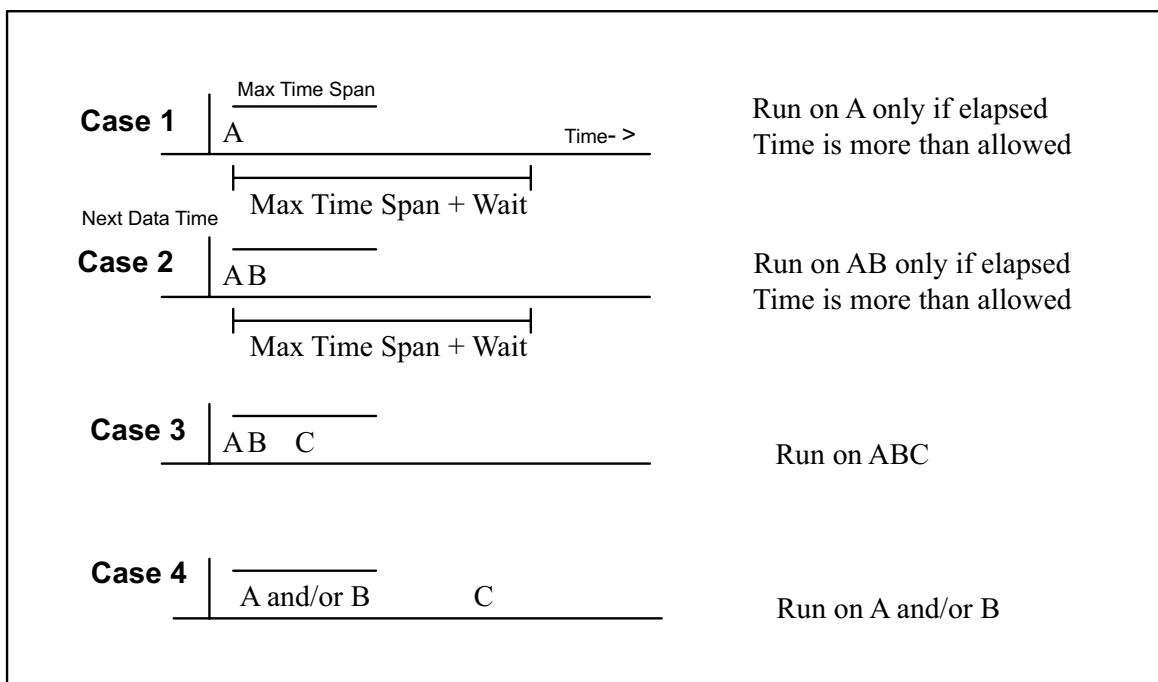
Figure 47   Composite Scheduling Example

The scheduling algorithm takes into account that not all radars are on the same schedule and not all radars work 100 % of the time.

Consider a simple case of products arriving from radars A, B, and C shown in the previous figure.

- A, B, and C refer to products from different radar sites.

- The vertical line at the left represents the **Product Scheduler** "Next Data Time". The scheduling algorithm only considers data ahead (to the right) of this line.
- The upper bar represents the "max time span" for making a composite. The max time span is set in the **COMP Product Configuration Menu**. See 5.2.3 Configuring COMP (page 138).
- The lower bar interval is the "elapsed time" between the time of arrival of a product and the current time.
- The products are positioned on the time line according to their "Data Time" (for example, start of a volume scan).
- Each site is ranked in priority for the purpose of assigning the "Data Time" of the composite. When the composite is made, it inherits the data time of the highest priority used site. The ranking is determined by the order of the site list in the **COMP Product Configuration Menu**.

Cases 1 and 2: Incomplete Data, Do Not Wait Forever

In these cases, the scheduler waits for the data to be completed until the time between the arrival time of the earliest product and the current time exceeds the time window plus the "Product arrival wait time" from **Setup**.

This provides tolerance for radar workstation clocks that are not exactly synchronized, or for variations in the network transfer time.

After running, the scheduler advances the "next data time" pointer to just after the data time of the latest product.

Note that only the radar computers must have their times synchronized. The system clock on the compositing computer is irrelevant.

Case 3: A, B, and C Are All Available and in the Time Window

Run on A, B, and C. This is the best case since all the data have arrived.

After running, the scheduler advances the "next data time" pointer to just after the data time of the latest product that just ran.

Case 4: Out of Sequence Data or Late Arrival

A and B both arrive and then C arrives out of the time window. Run on A and B only. After running, advance the next data time to just after the data time of the latest product that just ran.

A arrives then B or C arrive out of the time window. Run on A only. Advance the next data time to just after the data time of A.

Using Skip Times

The discussion above is for a schedule where the skip time is set to 0.

In this case, no data is ever reused, and the next data time is set to just after the last product used as inputs.

The only difference when using a skip time is that the next data time is set to the next multiple of the skip time after the first product that was composited. This guarantees that a composite is not skipped because of a single missing input file. It also means that in some cases the same input file may be used in 2 composites.

It is recommended that the **Max Time Span** be set to just under the skip time. Using a skip time is generally required if you are compositing radars with different times between volume scans.

5.3 **DWELL** Algorithm: Composite in Time

DWELL is comprised of a product and composites successive images of a product in time. Moving targets show a "streak" on the display.

DWELL is similar to the photographic technique of leaving the camera shutter open while photographing a moving target. The background is fixed and the moving target shows a "blur" behind it.

Figure 48   **DWELL** Schematic Example

The **DWELL** product algorithm is a "time composite" of the input products. Because of this, the product that is generated by the **DWELL** algorithm is of the same type as the input product. For example, if we "dwell" a **PPI** product, a new **PPI** product is generated. The choices for the type of data that are displayed in the **DWELL PPI** product are:

- *Original input data*
  For example, if a **PPI** of dBZ is used as input, then **DWELL** makes a time composite of dBZ.
  This is useful for showing where precipitation has fallen during the dwell time. When animated, the life cycle of growth, decay and motion of individual echoes can be easily seen.
- *Time*
  The data points in the **DWELL** product are coded by their age from the time of the most recent product in the **DWELL**. The display shows a color coded time history of echoes.

Because **DWELL** products show the integrated time history of echoes, they are a useful alternative to the **TRACK** product.

In airborne target detection, **DWELL** is used to track and automatically warn about the presence of aircraft or migrating bird flocks.

5.3.1 DWELL Algorithm Examples

DWELL is comprised of a product and composites successive images of a product in time. Moving targets show a "streak" on the display.

The example shown here is for **DWELL** applied to **PPI** products. The last (most recent) **PPI** product in the **DWELL** sequence is shown below. It shows a squall line in dBZ. The line was moving from East to West and developing/decaying as it traveled.
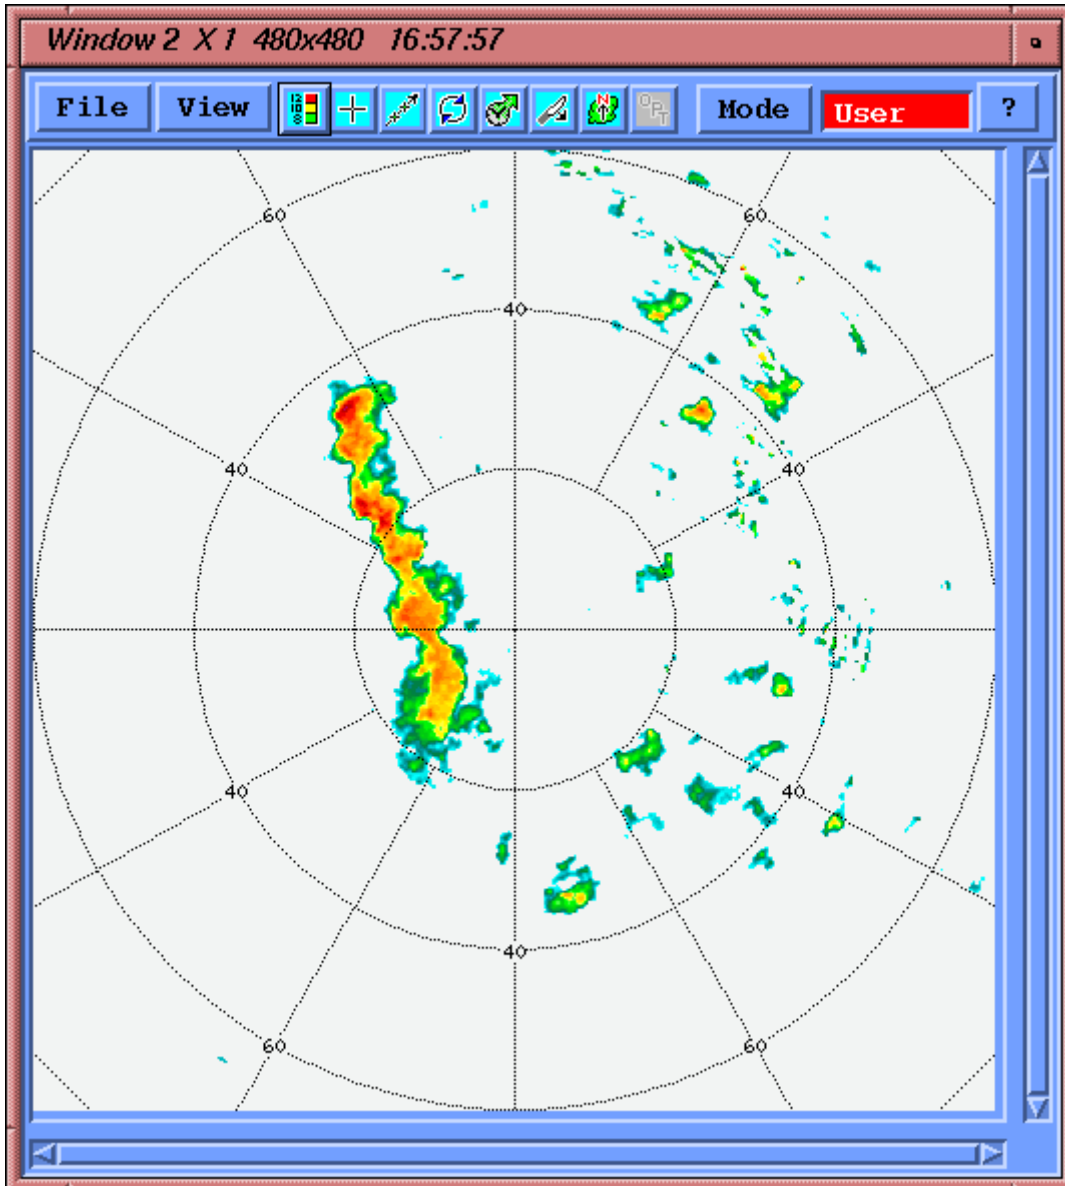
Figure 49   **PPI** of **dBZ** at a Single Time

The following figure below shows a **DWELL** of the **PPI** products for the preceding 2 hours. This is useful to show where precipitation has fallen during the **DWELL** period. The animation shows the expanding field of fallen precipitation and the echo motion.
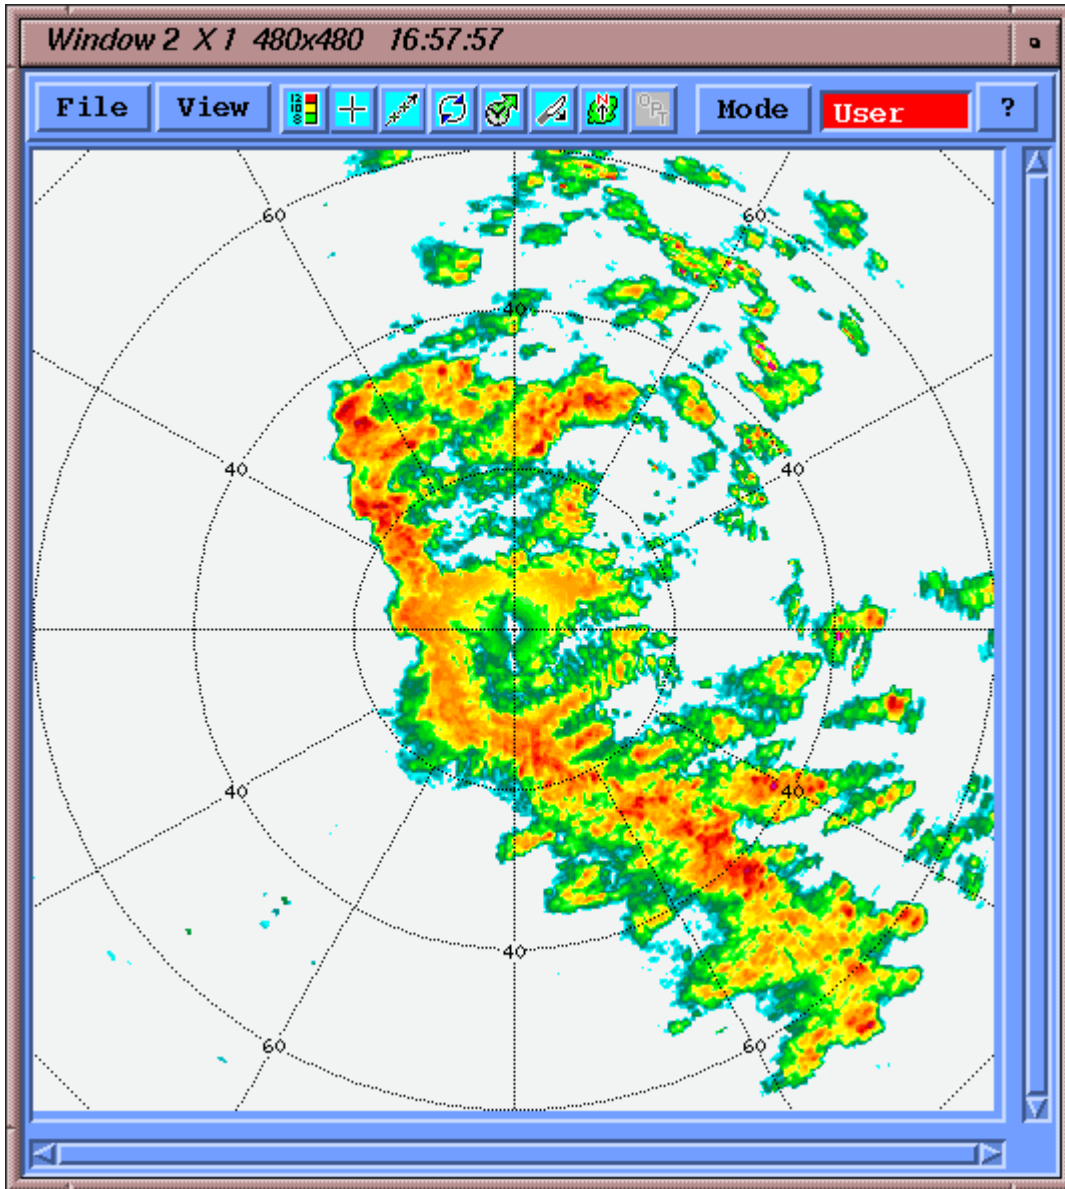
Figure 50   **DWELL** of **PPI** for the Previous 2 Hours Showing dBZ

The final figure in the example shows the age of the **DWELL** data points, that is, the number of minutes since the most recent **PPI** product in the sequence. The color scale is configured so that blue shows the most recent echoes while red shows the oldest echoes.

The example shows blue on the left (West) where the current echo is located. The red echoes on the right (East) are the older echoes from earlier in the sequence. This is to be expected for general East to West motion. Residual clutter shows a speckled pattern (for example, at 50 km (31.1 mi) to the Northeast).
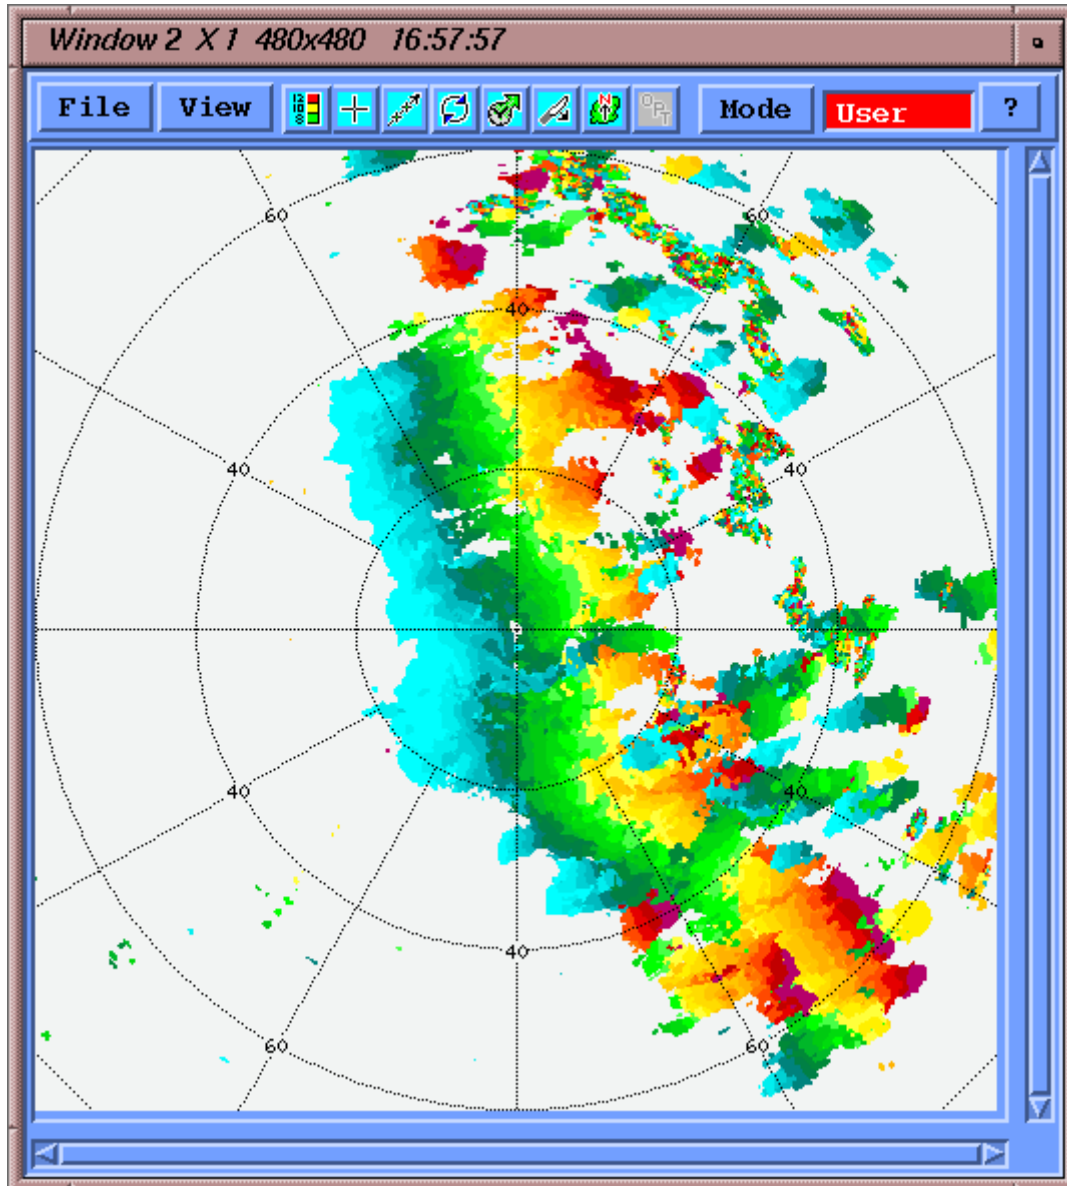
Figure 51   Dwell of PPI for the Previous 2 Hours Showing Time (Age)

5.3.2 Applying the DWELL Algorithm to Product Scheduling

The basic **DWELL** algorithm (without target warning) takes a series of products from the same site, of a given name and type and merges them together.

Product Scheduling

You schedule the **DWELL** algorithm in the **Product Scheduler**.

When the **Product Scheduler** detects an input product beyond the `Next Data Time`, this product and all products from the same site, within the specified dwell time (for example, the last 10 minutes), are combined to make a new product of the same type. The data stored in the product can be either the data values of the original product or the age of each point relative to the most recent product in the series.

> **i** If the **DWELL** product is based on a **PPI** product where the elevation selection in the **PPI** product configuration is set to the wild card *, then you must take care in scheduling to avoid excess CPU loading and performance degradation.
> To prevent the **DWELL** product being produced on the completion of every **PPI** sweep,Vaisala recommends that you set the **Product Scheduler > Skip Time** to the time required to complete a volume scan or longer.

Input Product Filters

You can apply the following input data filters to each input product before it is merged into the **DWELL**:

- *Min Filter*
  Discards any pixel in the input product with data value less than the specified minimum value. For velocity, absolute value is used.
  For the example of a **PPI** in dBZ, if the minimum value is set to 10 dBZ, all pixels with values less than this are discarded.
- *Max Filter*
  Discards any pixel in the input product with a data value less than the specified maximum value. For velocity, absolute value is used.
  For the example of a **PPI** in of dBZ, if the maximum value is set to 50 dBZ, all pixels in the input product with values greater than this are discarded.
- *Contrast Filter*
  Discards any pixel in the input product with data value not exceeding the corresponding data value in the computed background field plus a selectable threshold. For velocity, the absolute value of the difference between the background field and input product pixel is used. The background field is obtained by first creating a smoothed version of the **DWELL** product.
  For the example of a **VIL** product, if the contrast threshold is set to 3 (mm) (0.12 in), then a pixel in an input product must have a VIL 3 mm (0.12 in) greater than the corresponding pixel in the background field.

If a pixel fails any of these tests, then it is set to **thresholded**, that is, no data.

Data Merge

After filtering, each pixel in the input is "merged" into the corresponding pixel in the **DWELL** output. If, for example, there are 10 input products to merge, there can be up to 10 possible values for assigning the output value. The merge algorithm uses the maximum value. In the case of velocity, the absolute value is used.

Depending on the user selection, the algorithm then stores either the time of the merged pixel or its data value. Time is stored as the "age" of the pixel relative to the most recent input product in the **DWELL** sequence.

In the case of merging **WARN** products, the centroid statistics are merged together, that is, a new **WARN** product is created that combines all of the centroids in the time sequence of input **WARN** products.

Background Field

The contrast filter allows relatively strong echoes such as reflectivity cores, birds or aircraft to be isolated from relatively weak background echo. For the target warning algorithm, the contrast filter is essential to detecting airborne targets in weak clear air or weather echoes.

The background field is created by first running the **DWELL** algorithm without the contrast step and then smoothing the data with a 2D filter. Note that the Max and Min filters are used to filter the inputs to generate the background field.

If the input product is **PPI**s at different elevation angles, you can generate the background field separately for each **PPI** elevation angle instead of merging the **PPI**s regardless of angle.

The merge algorithm for the creation of the background field is the same as for the main **DWELL** algorithm (that is, maximum value), except that for velocity, the average value of the signed velocity is used. The reason for using the Max for everything except velocity is that weather echoes have a natural power fluctuation and these are better removed by the contrast filter if the background field is based on maximum, that is, the background field has larger values so that fewer weather speckles are passed. In the case of velocity the average value is used to characterize the background field since it does not make sense to use the maximum when every velocity is equally likely.

After the merge, a smoothing filter is applied. This is a 2D averaging filter set to 10 by 10 km (6.2 by 6.2 mi) or 60 × 60 pixels (whichever is less), centered at the smoother output pixel. A triangular weighting is used to assign more weight to the center pixels. The filter requires a selectable percentage coverage to produce an average value at the center pixel, else a value of "thresholded" is assigned to the center pixel. The threshold is typically set to 20 % coverage. This has the effect that isolated point targets such as birds or aircraft are excluded from the background field so that they are not filtered by the contrast filter.

An example of a smoothed contrast filed is shown below for the same case in the examples in .

Figure 52   Background Field for 2-Hour Dwell Showing dBZ

The smoothed background field has generally lower echo intensity as a result of the smoothing process. Note that there are no "speckles", that is, only large scale echo features remain. Thus the background field is a representation of the average background echoes. When this is subtracted from the input products and thresholded (in applying the contrast filter), the background echoes are removed, leaving only the high intensity cores or point targets such as birds or aircraft.

Diagnostic Output for Algorithm Verification

To assist with tuning the **DWELL** algorithm, especially in the case of target warning, the **DWELL** algorithm can produce diagnostic output.

To activate this output, select **Diagnostic** at the bottom of the **DWELL Configuration Menu**. Note that diagnostic output is produced even if you later disable the target warning, as long as the diagnostic button is enabled.

The diagnostic outputs are stored as type `USER` products. The naming convention is as follows:

- `USER: < productname >DAT or TIM`
  The main output of the **DWELL** product is a product of the same type as the input product (for example, a **PPI**) for either the input product data type (for example, dBZ) or time (age).
  The diagnostic `USER` product is for the data type that was not selected for the main output.
  For example, if **dBZ** is selected for the main output, the diagnostic `USER` product is for the time (`TIM` suffix). If time is selected for the main output, the diagnostic `USER` product is for the input data type (`DAT` suffix).
- `USER: < productname >CON`
  If the contrast filter is enabled, the diagnostic `USER` product with the suffix `CON` shows the background field. For an example, see .
- `USER: < productname >PIL`
  If target warning is enabled, the diagnostic `USER` product with the suffix `PIL` shows the "pile" product, that is, the number of target hits per pixel (X4) for the first in the sequence of the assumed motion vectors (that is, the first trial).
  For diagnostic work, the target Warning can be configured to run a single motion vector so that any speed and direction can be isolated for study.

5.3.3 Configuring DWELL Products with Target Detection Disabled
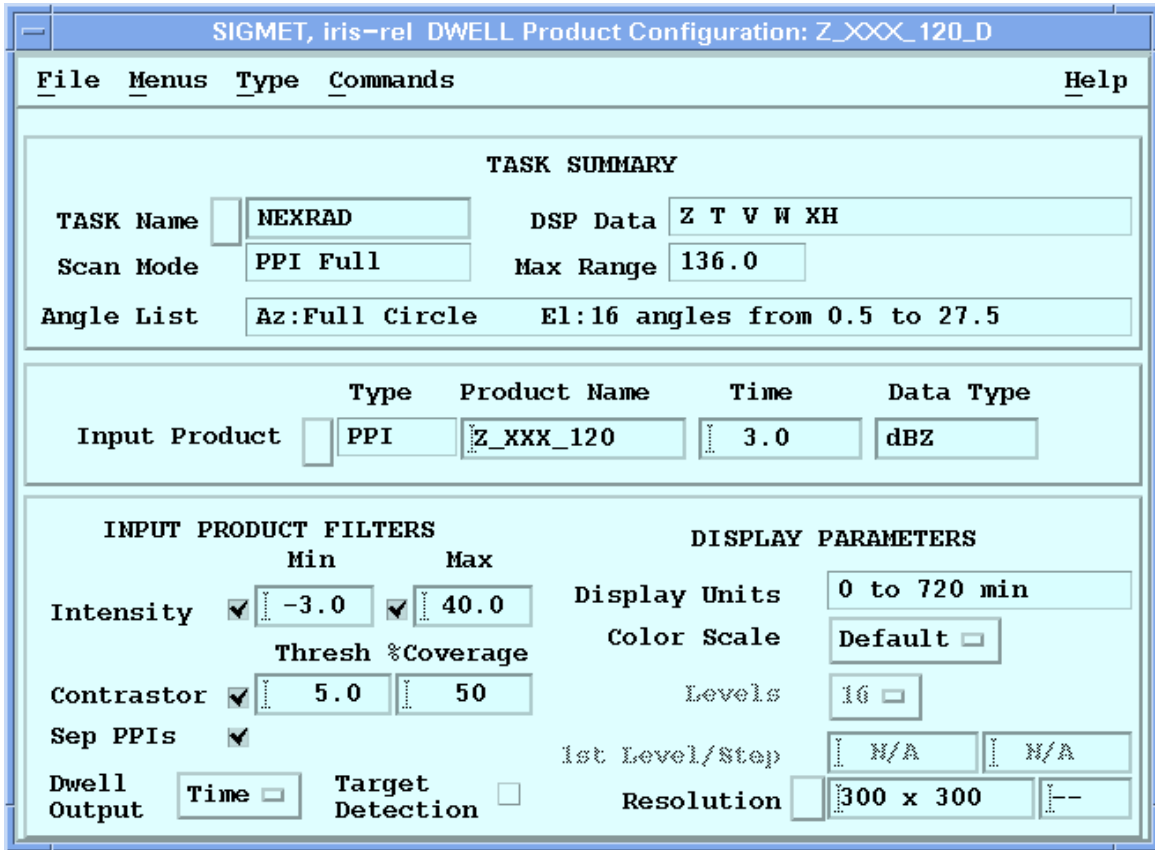
Figure 53   **DWELL** Example Configuration with **Target Detection** Disabled

▶ 1. Select **Type > DWELL**.

2. To name the **DWELL** product, select **File > Save As**.
   Vaisala recommends that the name of the **DWELL** product be the same as the input
   product with a D appended to the name with or without the number of minutes. For
   example:
   - Input product name: Z_005_250
   - **DWELL** product name: Z_005_250_D
   - **DWELL** product name: Z_005_250D10
     This indicates a 10-minute **DWELL**.

3. **Input Product Type** and **Name** to define the input product, that is, the product that you want to "dwell" together.
   You can select any of the following product types:
   - **BASE**
   - **CAPPI**
   - **HMAX**
   - **PPI**
   - **RAIN1**
   - **SRI**
   - **TOPS**
   - **VIL**
   - **WARN**

   When combining **WARN** products, the centroid characteristics of all the input products are combined together. Therefore, most of the menu features in the **DWELL** product are disabled, except the **DWELL** time selection.

4. Define the dwell **Time**.
   This is the time period in minutes for which products are dwelled together.
   For example, when a new **PPI** product is created or received, this and all **PPI** products within the last 10 minutes are included in the **DWELL PPI** product.

5. **DWELL** products using **PPI** as an input product, select the **Data Type**.
   If you select the **Dwell Output** to be **Data** (for example, in the lower left corner), the **DWELL** version of the **PPI** is made for this data type.
   In the example, a **DWELL PPI** of dBZ is generated.

6. To use the intensity filter, select **Input Product Filters > Intensity**.
   The intensity filter allows you to eliminate any data points that are less than the **Min** or greater than the **Max**.
   A useful application is to set a **Min** threshold to eliminate weak echoes so that the **DWELL** algorithm result shows only the strongest precipitation cores.

   > ℹ️ Check an example of the input data to determine how to best set the thresholds.

   a. Enable the filter by selecting the button.

   b. Type the value.
      The units are for the **Data Type** of the input product.
      In the example menu, this is dBZ.

7. To use the contrast filter, select **Input Product Filter > Contrastor**.

> **i** For target detection, the contrastor is necessary in order to pick-out an airborne target from weak clear air or weather echoes.

The contrast filter provides a way to eliminate relatively weak echoes so that only the embedded stronger echo cores remain.
It first computes a smoothed version of the **DWELL** (the background echo) and then subtracts this from the original input data to get the final "contrasted" version of the **DWELL**.
The **DWELL** algorithm runs twice: the first time to compute the background field and the second time during which the background field is used to threshold the original input data. For more information on the contrast algorithm see 5.2.3 Configuring COMP (page 138).
The parameters to configure the contrast filter are as follows:
- **Thresh**
  Each input product data pixel must be greater than the corresponding pixel in the average background field by this amount, else it is discarded.
- **% Coverage**
  The background field is produced by a smoother that averages data pixels in a 10 km by 10 km (6.2 mi by 6.2 mi) box that is passed over all the pixels in the first pass **DWELL** product. If the percentage of valid (non-thresholded) pixels in this box does not exceed the **% Coverage** limit, then a value of "thresholded" is assigned to this box, that is, there are not enough points to compute the average.

8. If your input product consists of **PPI** products at different elevation angles, you can select **Sep PPI's**.
For example, a **PPI** made with a * in the elevation angle filed has every angle in the volume scan. In this case, the background field can be in one of the following ways:
- Default method
  All of the angles are merged together to compute the background field.
- **Sep PPI's** enabled
  The background field is computed separately for each **PPI** elevation angle and the contrast thresholding is performed separately for each angle before merging the results.
  This provides better results when there are vertical gradients in velocity or reflectivity. If the vertical gradients are weak then it does not really matter which approach is used.

9. In **Dwell Output**, select one of the following:

    *   **Data**
        The **DWELL** algorithm produces a time composite product in the same units as the input products.
        In the example, a **PPI** of dBZ would be produced.
    *   **Time**
        The **DWELL** algorithm produces a product that shows the age of each point relative to the most recent input product.
        This display makes it easy to spot moving targets since they leave a streak of colors ranging from the oldest to the most recent.

    For examples, see 5.2.1 Composite Algorithm (page 136).

10. In **Display Parameters**, select the color scales.
    The available parameters depend on the value you selected for **Dwell Output** (**Data** or **Time**).
    You cannot change the output pixel numbers since it is controlled by the resolution of the input products.

More Information

‣ Configuring DWELL Products with Target Detection Enabled (page 162)

5.3.4 DWELL Target Detection Algorithm

The target detection algorithm uses a series of motion vectors, spanning the possible range of target speed and direction, to shift a time **DWELL** product.

Targets whose motion matches the assumed motion vector "pile". The depth of the pile provides a measure of the "coherent motion" of the target. Essentially this is an autocorrelation approach. The steps are:

1. A time **DWELL** product is made in the usual manner.
    Note that if the target warning is enabled, then even if a data **DWELL** product were requested, the time **DWELL** would be created for the warning.
2. The time **DWELL** points are shifted back in horizontal space according to a trial velocity vector.
    Targets that are moving at a constant velocity are "piled" on top of each other. The output of this step is an intermediate "pile" product which has the number of data points that are accumulated (piled) on each pixel.
3. A warning/centroid algorithm runs using a selectable threshold for the number of "pile" points.
    The approach is the same as in 3.25 WARN: Warning/Centroid Product (page 104).
    Note that the pile product units are 4X the actual number of points. For example, for a pile of 3 points, the representation would be 12 in the diagnostic pile product (PIL).
4. The previous 2 steps are repeated for each trial.
5. After all trials have run, the centroid products for each trial are combined to eliminate "double hits" that occur for the same flock. That is, there may be several trial winds that produce centroids of sufficient intensity (pile number).
    The combination combines all centroids that are within 3 km of each other by choosing the one with the largest pile number.
6. The combined warning product is checked for protected areas that have been hit.

7. The final output **WARN** product is named using the same name assigned to the **DWELL** product.

Make Dwell Time Product

t–2

t–1   Input Product Sequence

t0

Dwell Product at Time t0

Check Each Trial Motion Vector for Targets

t0   Shift by trial motion vector
t-2 t-1

Next Trial Motion Vector

3 hits
t–2
t-1
t0   Pile product

+B
Trial 1 WARN   Warning centroids based on the pile

Combine Warnings from All Trial Motion Vectors and Then Check Protected Areas

+B
Trial 1 WARN

+B
Trial 2 WARN   Final Target Warning

+B
Trial 3 WARN   +B
                              +B

+B   +B
Trial 4 WARN

Figure 54   Schematic of Target Warning Algorithm

Target Detection Data Format

The target detection algorithm generates an IRIS **WARN** product.

Part of that `warning_results` structure are 3 data values. For normal **WARN** products, these contain the average and maximum data values within the detected centroid. When generated with the target detection algorithm, these contain:

- The average and maximum height of the target returns (km).
- The average and maximum pile height.
- The average and maximum elevation angle of the target returns in degrees.

5.3.5 Target Detection: Input Task and Product Optimization

The airborne target detection, tracking and warning features of the **DWELL** algorithm are designed to detect aircraft or flocks of migrating birds and issue automatic warnings.

Aircraft and migratory bird flocks tend to move with a fairly constant speed and direction and this property is used to distinguish between airborne targets and other targets. In many cases, targets embedded in weak weather targets or clear air echoes can be distinguished.

As with any automatic detection algorithm, there is a direct tradeoff between the probability of detection and the false alarm rate. For best performance, you must optimize the scan task and the input product generation.

Scan Task Optimization

Target detection requires that extraneous targets such as clutter, second trip echo and noise be eliminated. The update rate should be fairly rapid to resolve the target motion and provide more continuous information on the target locations. Some knowledge of the likely altitudes of flight is required to construct the volume scan task.

Table 15   Optimizing Scan Task for **DWELL** Target Detection

Parameter	Description
Maximum range	60 ... 100 km (37.3 ... 62.1 mi).
	The maximum range of detection depends on the performance characteristics of the radar.
	Typically about 60 km (37.3 mi) detection range can be expected for a 1 to 1.5° beam width antenna.
Maximum resolution permitted by the signal processor	For example, 125 meters (410 ft 1 in).
	This depends on the processor performance characteristics.
Maximum PRF and minimum pulse width	--
2nd trip echo cancellation	On magnetron systems, SQI thresholding of the dBZ may be used instead of this.
	This discards non-coherent targets such as 2nd trip echoes.

Parameter	Description
Ground clutter filtering with an aggressive CSR threshold (lower value). alarms.	The value depends on the radar phase stability performance characteristics but is typically be in the range 10 ... 20 dB. The goal is to eliminate strong clutter targets which could produce false alarms.
SQI	Adjusted higher than usual. For example, approximately 0.45 This is to eliminate velocity speckles which could produce false alarms.
LOG	Adjusted higher than usual for weather. This is to eliminate speckles in intensity which could produce false alarms.
Speckle remover	Disabled so that single bin targets are passed. These could contain a point airborne target.
Elevation angles	Few elevation angles covering only the heights of interest to a range of approximately 60 km (37.3 mi). This allows more rapid updates. For example, for a 1° antenna beam, the 5 elevation angles of 0.5, 1.5, 2.5, 3.5, 4.5 would provide continuous vertical coverage up to 2 km (1.2 mi), from 25 ... 60 km (15.5 ... 37.3 mi). At 10 km (6.2 mi) the coverage would be to above the radar.
Scan rate	Moderately fast scan rate (>3 RPM) consistent with rapid updates, but still allowing approximately 50 pulses per antenna beam width. For example a PRF of 1000 Hz and a scan rate of 3 RPM (18 degrees per second) would provide 1000/18=55 pulses for each degree of antenna motion.

Input Product Selection

Table 16   Suggested Input Product Types for Target Warning

Product Type	Description
PPI dBZ or radial velocity at single or multiple sweeps.	Radial velocity is recommended for discriminating between airborne targets and background clear air or weak weather echoes. Airborne targets, such as aircraft or birds, move at a different velocity than the air so that, provided the targets are stronger targets, they produce a velocity anomaly which can be passed by the contrasting algorithm. If the motion of the target is perpendicular to the radar beam, the clutter filtering can remove targets so there is a possibility of some "blind" sectors because of this.

Product Type	Description
VIL Covering the layer of interest.	**VIL** is a useful product for target warning since it integrates echo over a selected layer. Single **PPI** products can be used, but since they are at only one elevation angle, not all heights are covered. However, a **PPI** product made with the wild card * in the elevation selection can be used to combine together all the **PPI** products in a volume regardless of the elevation angle. In this case, the **DWELL** algorithm runs at the completion of every sweep, and since the algorithm is intensive, the system may not have enough CPU power to do all of the **DWELL** products plus the other functions (for example, user display, animation, other product generation, communication). The **DWELL** product generation would then fall behind real time. Therefore if all of the **PPI** products are used, it is recommended that the **DWELL** product schedule entry use the **Skip Time** in **Product Scheduler**, set to the time between volume scans. Thus a new **DWELL** product is produced on the completion of the volume scan rather than on completion of each individual **PPI**.

Input Product Optimization

Table 17   Optimizing Input Products for **DWELL** Target Detection

Parameter	Description
Maximum range and the pixel resolution	Set so that the resolution of the original input data is preserved. For example, for 125-m (410 ft 1 in) bin spacing to 45 km (28.0 mi), the product pixel size should be set to 720 × 720 (corresponding to 90 km (55.9 mi) diameter and 0.125 m (5 in) per pixel).
Product smoother	Must be set to 0. This is so that isolated single-bin targets are not removed.

5.3.6 Configuring DWELL Products with Target Detection Enabled

Figure 55   Target Detection Configuration Example

> For information on **WARN** product options, see 3.25 WARN: Warning/Centroid Product (page 104).

▶ 1. In the **DWELL Configuration Menu**, select **Target Warning**.
   A new section of the menu appears at the bottom of the main menu (see above).

2. Configure the **DWELL** configuration options for target detection.
   Warning Symbol

   Enter the name of an icon to display. If no icon exists by this name, then the text that is entered shall be displayed at each centroid location. The icon automatically rotates to point in the direction of target motion.
   In the example above, **Plane** is an icon of an airplane that is flying North. Bird icons are also provided.

   Correlation Threshold

   This is the threshold for the number of points in the "pile" product X4. For example, the number 8 represent 2 points. Note that the smoothing in the pile product typically reduces the number of points from the maximum value. Thus a pile product value of 2 would correspond to a greater number at very sharp maximum.

   > The correlation threshold is the primary tunable parameter to optimize the probability of detection and false alarm rate.

   > The appropriate value of the correlation threshold depends on the length of the **DWELL** sequence. Longer **DWELL** times (more input products) produce larger "piles".

   Add Wind, **Use VVP**, **Speed**, and **Heading Range**

   The track of the targets is used during the "pile" process to add up targets that are moving at the expected velocity. The velocity of the targets is the result of the target's flight (air velocity) and the wind. These fields are used to account for these effects.
   Add Wind specifies whether you want the wind effect to be added to the target flight velocity vector. If this is enabled, the **Speed Range** is the expected range of speed of the targets in still air, that is, their flight speed.
   If you select **Add Wind**, select **Use VVP** to control where the algorithm gets the wind. The **UNFOLD VVP** product is used for mean wind information.
   Otherwise the **Setup** product section has a default wind speed and direction. In tropical maritime trade wind regions this could be set fairly reliably to a single value of speed and direction. For most regions, the default wind is not usually a very good value. You can change this value in real time using the **setup_change** utility.

   > If you choose not to use the **Add Wind** feature, then the range of speeds must include the expected variation in the wind speed as well as the target speed. The greater range of speeds results in more false alarms and requires more trial winds (slower algorithm performance).

   Target Size

If enabled, the **DWELL** "pile" product is further qualified to remove all large piles to reduce false positives.

If there is a squall line moving in the same direction as the expected airborne targets, it usually generates target warnings. These warnings generally have a larger area than the true targets, so you can filter them out.

Type the number in square km.

Diagnostic Output

If enabled, **DWELL** produces three diagnostic products. These are stored as `USER` products under the same product name assigned to the **DWELL** product and a suffix indicating the type of diagnostic (`DAT` or `TIM`, `CON`, and `PIL`).

See 5.3.2 Applying the DWELL Algorithm to Product Scheduling (page 149).

More Information

‣ Configuring DWELL Products with Target Detection Disabled (page 154)

5.3.7 Target Detection: Migratory Bird Examples

The following examples are for bird tracking during a fall migration. Targets are primarily flocks of large storks comprised of 30 or more individuals traveling from Northeast to Southwest (heading range of 200 to 250) at about 12 m/s (39 ft 4 in / s) (speed range of 8 to 16 m/s (26 ft 3 in to 52 ft 6 in / s)).

These figures show products and displays produced by the Target algorithm.

Input PPI's of Radial Velocity

The input data were based on a volume scan of four elevation angles at 0.5, 1.0, 2.0, and 3.0° elevation angle. The input resolution is 125 m (410 ft 1 in).

PPI products of radial velocity to 45 km (28.0 mi) range for all angles are used, that is, a product named **V_XXX_045** is generated with the wild card * in the elevation selection. A product pixel resolution of 720 x 720 is used to match the 125-m (410 ft 1 in) input data resolution.

The following figure shows an example of a 3° elevation **PPI**. There is a clear air echo only and winds are generally weak. At 15 km (9.3 mi) South of the radar there is a velocity anomaly that can be later identified as a bird flock. There are also other airborne targets, but since they may only be 1 or 2 pixels, it is difficult to separate these from the background noise and clear air echo.

Background Field
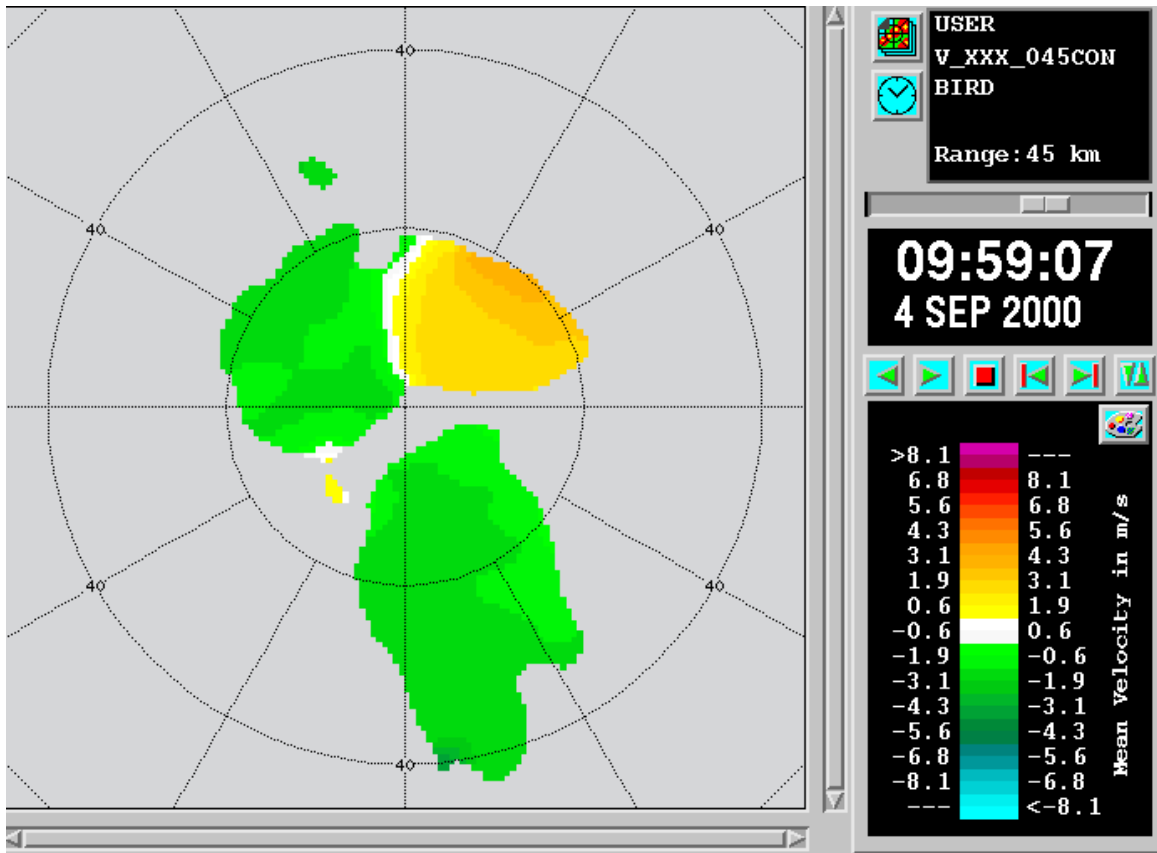
The contrast field is a smoothed version of the **DWELL** data which is used as a dynamic threshold for the contrast filter. For this example a minimum coverage criterion of 25 % was used. The **DWELL** is 10 minutes. Note that aggressive clutter filtering has removed nearly all of the clear air echo at 0 velocity.

Note that only the large scale echoes are included in the background field. This is important, otherwise airborne targets would be removed during the contrast step when this background velocity is subtracted.

Data Dwell Product
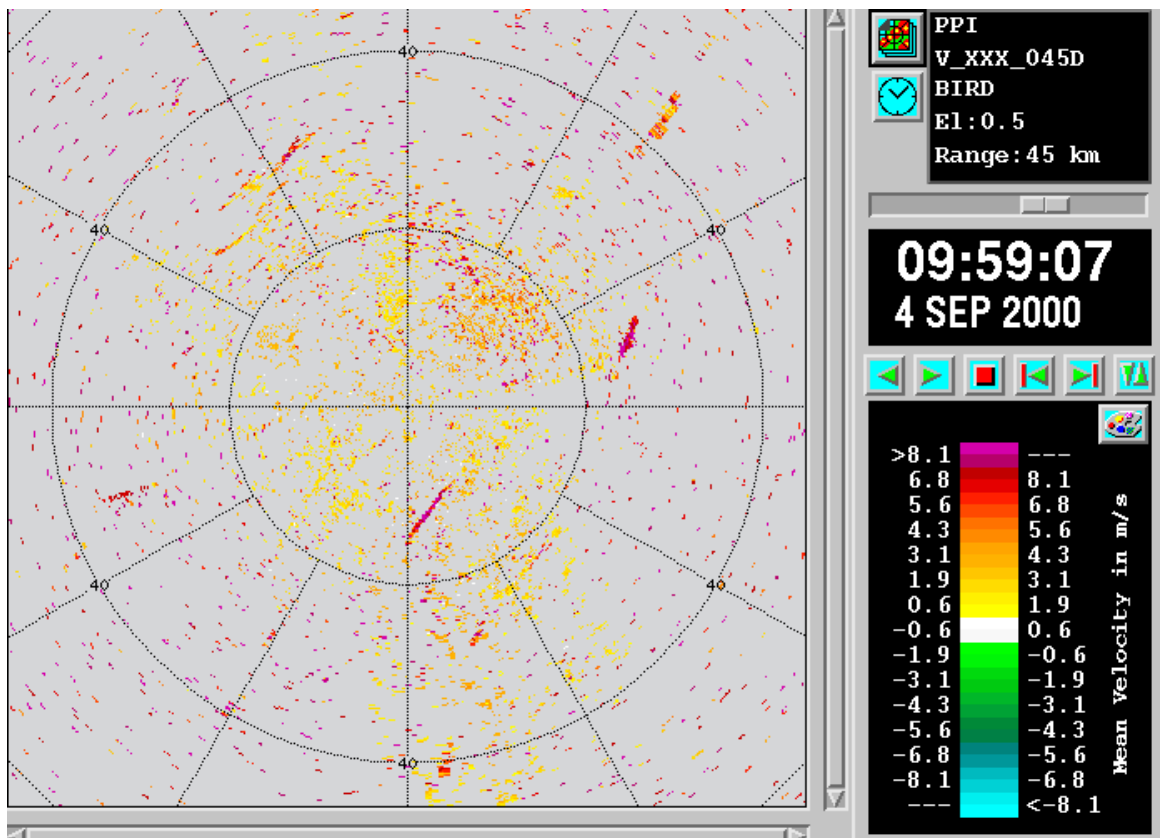
The data **DWELL** product (data was selected as the primary output) is a **PPI** product combining all of the **PPI** products from all elevation angles for the prior 10 minutes (22 **PPI**s in this case). Only positive values are included since the output is the absolute value of the velocity anomaly as compared to the background field. A contrast threshold of 1 m/s (3 ft 3 in / s) was used, that is, only data points that differ from the background field by more than 1 m/s (3 ft 3 in / s) are included in the **DWELL PPI** product.

This shows streaks of velocity anomaly which are likely bird targets. Note that all of the speckles in all of the **PPI** products are unfortunately passed through the **DWELL** algorithm. Also note that most of the clear air echo has been effectively removed by the contrast filter.

Time Dwell Product

The time **DWELL** product shows the age of the points in minutes. The most recent points are represented in blue and the older points in red. This clearly shows the streaks of motion from northeast to southwest. This is even more apparent in animation since the streaks move across the screen to the Southwest.

The random noise shows a chaotic speckled pattern in time. There is a streak of interference echo which shows only a single color (to the southwest at 30 km (18.6 mi)).

Pile Product (2X Zoom example)
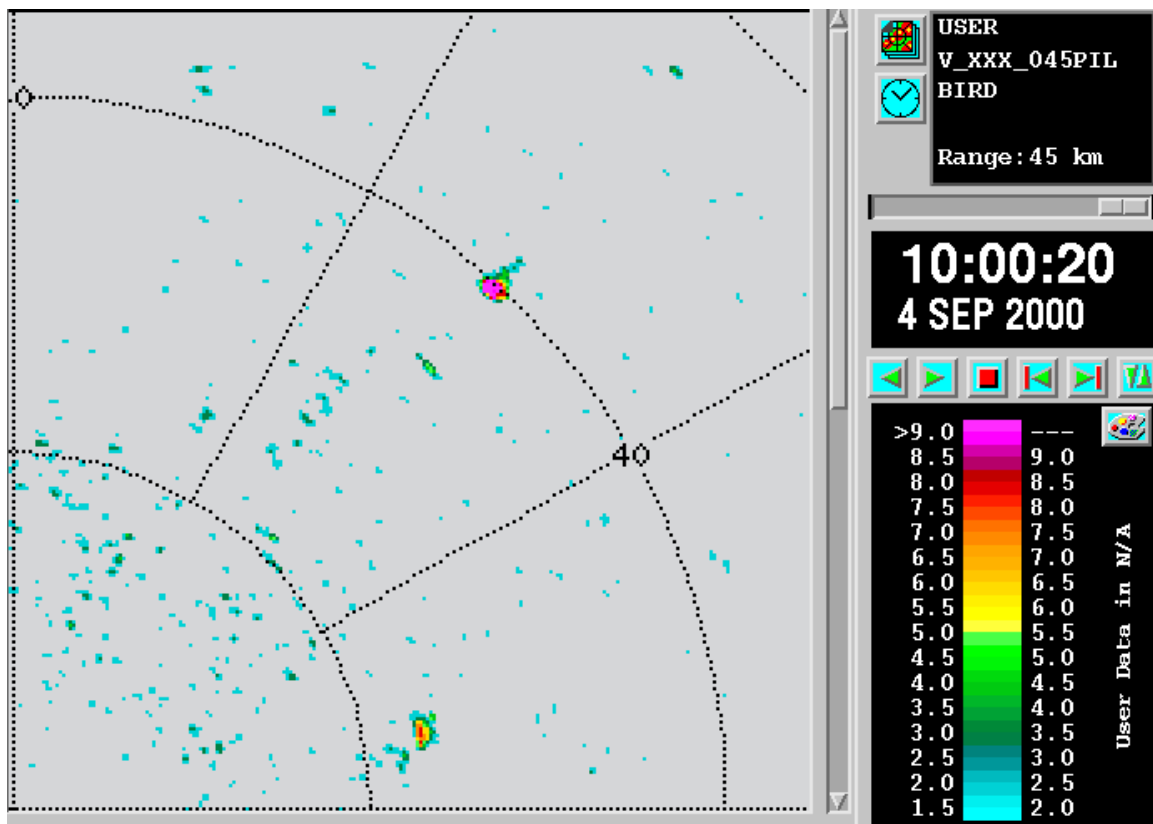
This shows the result of correlation step in which the time **DWELL** points are "piled". Since each trial wind produces a different pile product, the diagnostic output only shows the pile product for the first trial wind. In this case the result for the particular trial wind of 210° and 12 m/s (39 ft 4 in) was selected by inputting only these values into the limits for the speed and direction ranges.

The results show that the coherent streaks of echo as seen in the **DWELL** time display are effectively "piled-up" by this trial wind, while background noise produces only weak values. Not all streaks produce a large value since the targets may be moving at a different velocity. This is why it is important to use a spectrum of trial winds.

WARN Product

The final result is a warning product based on a composite of all warnings for all trial winds. This example used a threshold of 8 for the correlation threshold (pile product value). The icon here is drawn to make a clear indication of the type of warning.

Alternatively, the warning display can be configured via the output options button to display the numerical value of the height in Kfeet of the bird flock beside the icon, or to display the speed in knots of the flock. The height information is only available for track detection based on **PPI** input products since there is no height information available from the other input products that can be used for **DWELL** (for example, **VIL** and **TOPS**).

When this is animated, the birds appear to fly across the screen at their proper velocity. Some icons may drop in or out if they are small flocks on the edge of detection. Occasional false alarms show as isolated flashes during animation. The primary tuning parameter to balance the probability of detection and false alarm rate is the Correlation Threshold (pile height).

Dwell WARN Product

The example shows the results of the **DWELL** algorithm run on **WARN** products generated by the bird warning feature. This display is recommended for use by air traffic controllers who must advise pilots of possible bird activity. In the example above, the individual bird warnings are dwelled for 45 minutes. The display shows the selected icon (in this case a large bird) at the position of the most recent warning in the dwell sequence. Prior warnings are shown using a different icon which is also rotated to the direction of motion. The prior warning icon name is made by appending "d" (short for dwelled) to the primary icon. In this case we get *birdd.xbm*.

The result is a track style display that shows the bird activity over the dwell time. Inactive tracks do not have the large bird icon at the end. Tracks can become inactive if the birds fly into an obscured region or if the flock is small near the threshold of detection or if the birds stop to roost or perhaps "funnel" in a local thermal to gain altitude. However, even inactive tracks show controllers that there is confirmed bird activity in the area near the last sighting.

5.4 **GAGE**: Raingage Reports

The **GAGE** product stores reports from rain gage sensors for display and use in product generation. These products are produced using an input pipe from data sent by a weather station system.

Because **GAGE** products are not produced from radar data, they do not need to be configured or scheduled.

Each **GAGE** product file contains reports from all the rain gages, locations, and hourly precipitation amounts for one time interval. Usually the time interval is one hour, but it can be, for example, 15 minutes.

The display shows a color-coded **GAGE**icon (user definable) of precipitation. Clicking on the **GAGE** icon provides a time history of **GAGE** amounts for a selectable number of hours.

More Information

‣ Hydromet Raingage Correction Overview (page 377)
‣ Product Parameters (page 25)
‣ RAIN1 Scheduling with GAGE Product (page 386)

5.4.1 Displaying GAGE
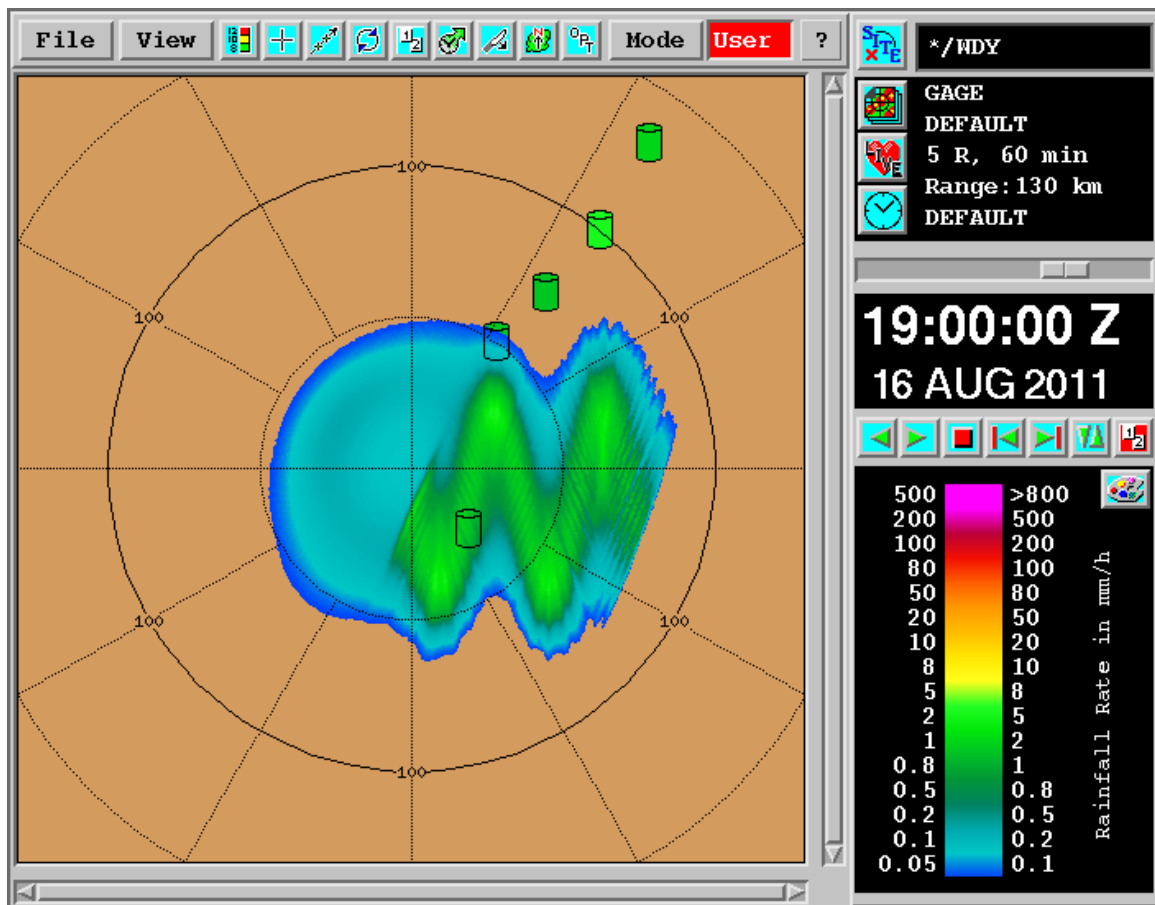
Figure 56   Raingage Product Display

You can display **GAGE** products in the **Quick Look Window**, as well as in other display outputs.

GAGE products are displayed in the form of a situation display showing a graphic of a rain gage can at the location of each gage.

Each gage is colored with the color for the rainfall accumulation measured by the gage. See the previous figure for an example of this display with radar data underneath.

The product-specific output option for gage products allows you to specify if the gage label is shown next to each gage.

▶   1. In the **Quick Look Window**, display a **GAGE** product.

2. Place the cursor in the display window and right-click the mouse.
   IRIS checks for the nearest rain gage location, and shows a rain gage datagram with 12 hours of history.
   The following example shows the hourly rainfall accumulation for the last 12 hours.
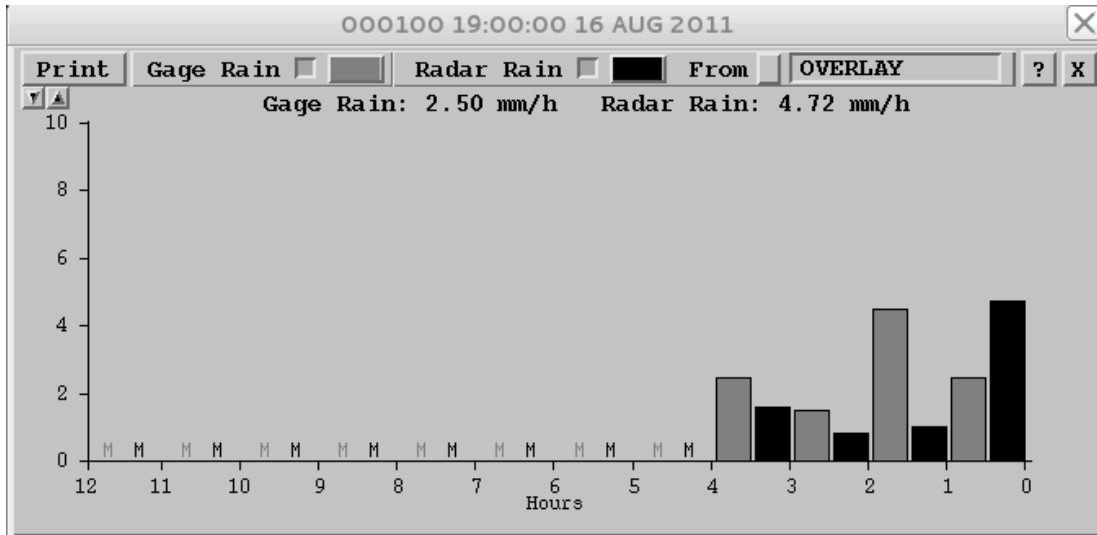
Figure 57   Raingage Datagram

You can configure the graph show just the rain accumulations from the rain gage, or radar rain accumulations from **RAIN1** products, or both.

3. If you are displaying radar rain accumulations, select which named **RAIN1** product to display from.
   The selected **RAIN1** product becomes the default **RAIN1** used as the topmost overlaid product. If there is no overlaid **RAIN1**, then its default is off.
   If you change the default **RAIN1**, future datagrams use the new product, until all datagrams are dismissed.
   You can show up to 10 rain gage datagrams simultaneously. This allows you to compare different gages, or **RAIN1** products.

5.5 **NDOP** : Multiple Doppler

NDOP provides a mechanism for including velocity (Doppler) data from another radar site now enabling the measurement of horizontal wind in addition to vertical wind.

A single Doppler radar measures one component of a 3D wind vector, that is, the component of wind towards or away from the radar: the radial wind. We can think of the 3D wind vector as being made up of the 2D horizontal wind and the vertical wind. Typically the horizontal wind is much stronger than the vertical wind, except in local areas of strong convection.

With a second Doppler radar, we can measure 2 components of the wind, provided that the radar beams are crossing at some angle (not parallel). If the vertical component of the wind is assumed to be small, then with 2 radars we can measure the horizontal wind.

The following figure shows a plan view (looking down) of the dual Doppler geometry. The horizontal wind is sampled by each radar which measures the radial component of the wind, `Vr1` and `Vr2`. These two radial velocity measurements at the same point, along with some simple trigonometry, can be used to estimate the horizontal wind at the point, provided that the crossing angle is not 0° or 180° (parallel beams).

Dual Doppler Geometry Plan View

Figure 58   Dual Doppler Geometry Plan View

This means that along the baseline (crossing angle 0°) or at far ranges (crossing angle 180°) we cannot use the algorithm, since the two radars essentially measure the same radial wind.

While in theory the crossing angle could be anywhere between 0° and 180°, our radial velocity measurements are not perfect so that in practice, a crossing angle of <20° leads to unreliable wind estimates.

The **NDOP** product performs the dual Doppler algorithm to obtain estimates of the horizontal wind field. **NDOP** makes the following assumptions:

- The vertical wind is assumed to be weak as compared to the horizontal wind.
- The input radial velocity measurements represent the radial velocity of the air, that is, the velocity is unfolded and corrected for particle fallspeed.

The output of the **NDOP** product is a multi-level grid of horizontal wind vectors. These can be displayed by themselves, or overlaid on other echoes.

More Information

‣ NDOP Output Options (page 256)

5.5.1 Input Velocity Corrections

Corrected Radial Velocity `Vc`

IRIS corrected velocity (`Vc`) can include corrections for fallspeed and velocity folding. The fallspeed correction is based on `VT–Z` relationships above and below the melting level. The unfolding correction is based on a **VVP** product.

`Vc` can be generated either:

- When the data are collected from the signal processor.
  This is configured in the **TASK Configuration Menu**.
  This is recommended for real time operation.
- When RAW restored from tape or received over the network are re-ingested (to make ingest files).
- This is configured in the **Setup** utility.
  Note that if `Vc` is already in the data when they are re-ingested, it is re-calculated from `V`. When `Vc` is generated, the uncorrected radial velocity `V` is still preserved.
  This is recommended for archive data or in systems where the communication bandwidth is limited so that the extra burden of transmitting `Vc` (as well as the standard radial velocity) is too much for the network.

Fallspeed Correction for `Vc`

A key assumption is that the vertical airmotions are weak as compared to the horizontal airmotions. This means that the radial winds are assumed to be caused by the horizontal wind only. However, while vertical airmotions may be weak, the fallspeeds of the hydrometeors (of order 1 to 10 m/s (2 to 20 knots) for rain) can make a significant contribution to the radial velocity. Therefore it is necessary to correct the radial velocities for the effect of fallspeeds.

The effect of particle fallspeed depends on the sine of the elevation angle. For example at 0° elevation, the fallspeeds do not affect the radial velocity. At 30° elevation angle (a typical maximum elevation in a volume scan), then half of the fallspeed would be observed (sine 30 = 0.5). Thus a 10 m/s (20 knots) fallspeed (hail and rain mixed) would contribute 5 m/s (10 knots) to the radial wind which is significant.

The fallspeed correction in `Vc` is made using a `VT–Z` relationship (terminal fallspeed - reflectivity). These take the general form of $VT=aZ^b$.

Since the particles are very different above and below the melting level, it is important to use different `VT–Z` relationships for these 2 cases. The default relationships used in IRIS are:

- Above the melting level (snow and graupel) $VT = 0.8\ Z^{0.06}$
- Below the melting level (rain) $VT = 2.70\ Z^{0.11}$

Here `Z` is in mm^6/m^3 and `VT` is in m/s.[1]

These relationships are entered by your system manager in the **Setup** utility. In addition, the average height of the freezing level is input for each month of the year.

---

[1] *For a discussion of VT–Z relationships refer for example to a text on radar meteorology such as Battan, Louis, J., 1973: Radar Observation of the Atmosphere, University of Chicago Press, p 132.*

The `VT–Z` relationships, and freezing level estimation are not perfect, but they can improve the radial velocity estimates compared to performing no correction. The corrections obtained are 1 … 5 m/s (2 … 10 knots), depending on the elevation angle. The corrections have little effect (<1 m/s (<2 knots) correction) for elevation angles less than 5°.

Radial Velocity Input for NDOP- Unfolding for `Vc`

The **NDOP** product assumes that velocities are unfolded or dealiased. A Doppler radar has a limit on the unambiguous velocity which is:

```
Vu = +- (Wavelength*PRF)/4
```

Radial velocities that exceed this are said to be folded. On a color display with blue representing radial velocity toward the radar and red representing radial velocity away from the radar, a fold appears as an adjacent blue-to-red color shift.

Unfolding can be performed in the signal processor itself (for example, using the dual PRF technique), or can be done by IRIS when `Vc` is generated.

The unfolding for `Vc` is based on a **VVP** product which assumes that the wind field varies linearly with distance from the radar. The **VVP** product used by `Vc` must have the special name `UNFOLD` and must be available on the system where the unfolding is being performed. This technique achieves 3 times unfolding. For example, for an S band system operating at 1000 Hz PRF (range 150 km (93.2 mi)), the unambiguous velocity is ±25 m/s (±50 knots). With unfolding, the unambiguous velocity is ±75 m/s (±150 knots), which can handle most extreme meteorological situations.

As with the fallspeed correction, the `Vc` unfolding correction can be made at ingest when data are collected from the signal processor (selected in the **TASK Configuration Menu**) or at reingest (selected in **Setup**).

The **VVP** unfolding technique works well if there is adequate echo coverage and that the wind field is approximately linear in its variations. In the vicinity of strong shear lines or fronts, this may not be the case.

> If you are using dual-PRF velocity unfolding and it is adequate to prevent folding it is recommended that you not use the **VVP**-based unfolding in `Vc` since you may harm the data in some extreme cases.

5.5.2 Configuring NDOP Products
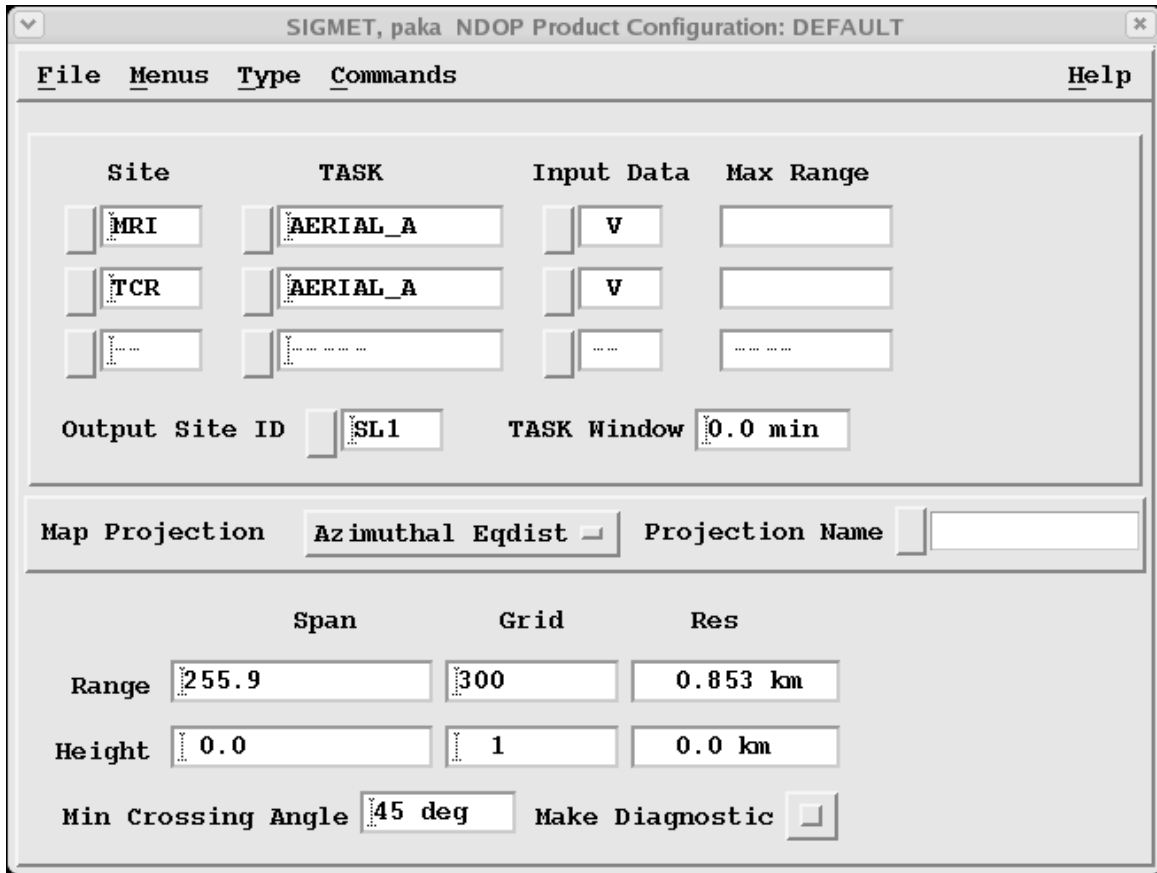
Figure 59   **NDOP** Example Configuration

1. In the top section of the **NDOP** configuration window, specify the tasks and sites.

   a. To select a site, select the **Site** button.

   b. To select a task based on the ingest files on your system, select the **TASK** button.
   If the task you want is not in the list, type the task name. You can use IRIS wildcard conventions, such as:
      - PPI*
      Use any task that starts with **PPI**.
      - PPIVOL_A
      Use only the **A** part of the hybrid task.
      - PPIVOL_BC
      Use only the **BC** parts of the hybrid task.

      If you have specified a single task, the range field is filled-in with the maximum range of the data for that task.
      If you use a wildcard or hybrid task specification the **Range** field is blank.

   c. Select the input data, V or Vc (radial velocity or corrected radial velocity).

   d. To enter a site ID associated with the product output, select **Output Site ID**.
   When the **NDOP** products appear in inventories such as the **Product Output Menu** or **Quick Look Window**, they are associated with this site.

   e. In the **TASK (time) Window**, specify is the maximum time difference that is tolerated for combining the data from the two radars.
   If the difference in the task start times is greater than this value, **NDOP** does not make the product. This prevents problems caused by advection and temporal changes that could cause inappropriate winds to be combined. A value of 2 minutes is typical.

2. In the middle section of the **NDOP** configuration window, specify the map projection for the product.
   See 3.2.2 Configuring Map Projections (page 21).
   The projection should be at a convenient point perhaps between the radars and cover a range that is no more than approximately 4 times the baseline length. This is typically the useful range of the **NDOP** product.
   If you do not specify a **Projection Name**, the data is centered on the location of the first radar site in the list.

3. In the bottom section of the **NDOP** configuration window, specify the output resolution by specifying the characteristics of the grid.
   The **NDOP** product produces wind vectors on a 3D grid.

   a. Specify the **Range Span**.
      The **Range Span** is the distance in km from the center of the projection, east to the edge of the projection, analogous to a radar range.
      It is specified in the named projection. If you are not using a **Projection Name**, type the range.
      A max range of about 4 times the baseline length is typical.

   b. Specify the **Range Grid**.
      The **Range Grid** is an integer that specifies the number of points in the output array representing the **Range Span**.

   c. For the **Height Span**, specify a maximum and minimum height (km).

   d. Specify the **Height Grid**.
      The resolution is then calculated from:

   ```
   Height Res = (Max Height – Min Height) / (Height Grid – 1)
   ```

   The Resolution (**Res**) is a display-only field calculated from:

   ```
   Range Res = (Range Span) / (Range Grid)
   ```

   In the example, the **Range Span** is 40 and the grid is 16, therefore the resolution is 2.5 km (1.6 mi). The actual number of points in the output array that is made by **NDOP** depends on the aspect ratio of the projection, which may not be square.
   For a square projection and a grid specification of 16, the output array would be 32 × 32.

4. In the bottom section of the **NDOP** configuration window, specify **Min Crossing Angle**. This is the minimum crossing angle that is used. The recommended value is 20°. The same value is used by the algorithm to specify the maximum crossing angle, that is, the `max crossing angle = 180 – Min Crossing Angle`.

5. If needed, enable **Make Diagnostic**.
   When enabled, the **NDOP** product makes additional velocity **CAPPI** products, with product names `NDOP_CAPPI_#`. These are the input data used by **NDOP**.

5.5.3 NDOP Display and Algorithm

The following example is an actual **NDOP** display showing the low-level (1 km) winds for a projection 40 km (24.9 mi) across using two S-band radar systems. [2]

---

2)   *Courtesy of the Hong Kong Observatory*

The meteorological situation is an event of strong Northerly winter monsoon flow over China colliding with Easterlies over the sea. The overlay map shows the contours of the mountainous terrain North of Hong Kong. Note the 50 knot winds funnelling through the mountain pass to the North of Hong Kong.

The "+" points in the display show the regions where the crossing angle criterion was met. These regions are on either side of the baseline that connects the 2 radars. However, if there are no weather targets, a wind cannot be calculated.

Figure 60   Example NDOP Display

The algorithm takes the following steps to compute the grid point winds:

1. Create 3D **CAPPI** products of radial velocity for the 2 radars in the common projection at the grid height spacing.
   The horizontal resolution of the **CAPPI** products is subject to the following constraints:
   - The resolution of the **CAPPI** pixels is set to be twice the spacing of the input radial velocity data. For example, for 125-m (410 ft i in) input bins, the **CAPPI** resolution is set to 250 m (820 ft 3 in).
   - The number of **CAPPI** pixels for each output resolution grid element should be at least 9 (3 × 3).
     If this is not the case, then the resolution of the **CAPPI** is increased.
     In the same example, if the output resolution were 2 km (1.2 mi) then per the 250-m (820 ft 3 in) pixel spacing in (1) there would be 8 × 8 = 32 pixels for each output grid point.
   - The maximum number of pixels in the **CAPPI** is 1100 × 1100. The **CAPPI** is clipped at this value.
     For a 40 km (24.9 mi) range span (80 km (49.7 mi) total across the output array), there would be (4 pixels/km)*80 (km) = 320 pixels in the CAPPI (320 × 320 for a square projection).
2. The data from the radial velocity **CAPPI** products are processed with the multiple Doppler algorithms to obtain a grid of (x,y) wind vectors at the original **CAPPI** resolution (for example, 320 × 320).
3. The high-density wind vectors are averaged to reduce the data to the final output grid. Continuing the example, with 4 pixels per km in the **CAPPI** and a an output resolution of 2 km (1.2 mi), we would average 8 × 8 = 32 wind vectors for each output grid point.
4. To eliminate noise and speckle effects, a grid point is thresholded if there are fewer than 3 values to average or less than 25 % coverage.
   In the example, the 25 % would correspond to 8 wind estimates.

When the vector averages are computed, a Wind Quality Index (WQI) is also calculated and stored with the data. This is used for thresholding when the product is displayed. For information on adjusting the threshold, see .

The WQI is computed from the variances $\#\texttt{sigma}\#^2{}_x$ and $\ddot{}\,\#\texttt{sigma}\#^2{}_y$ of the individual components of each (x,y) wind vector that is computed from the **CAPPI** data. For every **CAPPI** velocity pair that contributes to the final X and Y wind, the 2 variance terms are also computed so that WQI can be derived as:

$$WQI = 1.0 - \left(\sqrt{\frac{\sigma_x^2 + \sigma_y^2}{2}} / V_{norm} \right)$$

where V_{norm} V_norm is a normalization term, and is the standard deviation that would result if uniformly distributed random vectors at half the Nyquist velocity were input to the algorithm, that is, that case would produce a WQI of 0.

$$V_{norm} = \left(\frac{V1_u + V2_u}{2} \right) / 2\sqrt{2}$$

where $V1_u$ and $V2_u$ are the Nyquist velocities for the two **CAPPI** input products.

> ℹ️ WQI is set to 0 if the above calculation yields a negative number. Zero corresponds to a terrible fit among the dual Doppler winds that are averaged into each output grid; and 1.0 corresponds to perfect agreement of all the data.

5.6 **SHEAR**: Wind Shear

Figure 61   **SHEAR** Example Configuration

SHEAR detects wind shear in the atmosphere, allowing the detection of microbursts, gust fronts, mesocyclones, cold fronts, and atmospheric waves.

Wind shear is associated with many phenomena:

Microbursts
> Associated with convective storms. Extremely hazardous to aircraft during landing or takeoff. Microbursts are characterized by positive values of the radial shear (strongly divergent outflow) in a roughly circular region, typically less than 3 km (1.9 mi) in size.

Gust Fronts
> Caused by cold outflow from a convective storm (perhaps a microburst) colliding with the surrounding air. They are characterized primarily by negative values of the radial shear (convergence). However, depending on the geometry, they can also create positive values of the radial shear and azimuthal shear of either sign.

Mesocyclones
> Characterized by rotation. Mesocyclones are associated with tornadoes. The azimuth shear is used to detect mesocyclones.

Cold Fronts
> Similar to gust fronts, but much larger in extent.

Atmospheric Waves
> Produced at a variety of wavelengths and intensities. The shear values can be positive or negative, depending on the nature of the wave and the "phase" being observed.

IRIS can compute the following types of basic shear values for the radial component of the wind:

- Radial shear
  Computed by differencing the radial velocity in range. Positive values are for radial velocity increasing (more away) with range. Divergence of the radial wind is positive.
- Azimuthal shear
  Computed by differencing the radial velocity in azimuth. Positive values are for radial velocity increasing (more away) in the clockwise direction. This corresponds to positive vorticity.
- Elevation shear
  Computed by differencing the radial velocity in elevation. Positive values are for radial wind increasing (more away) with height.

In addition to the basic shears, combinations are computed using RMS values. For example, the total combined shear magnitude is:

$$SQRT\left(RAD^2 + AZ^2 + EL^2\right)$$

Where RAD, AZ, and EL denote the basic shear values. Combined shears are positive.

Figure 62   Schematic Examples of Wind Shear

> ℹ️ Vaisala does not warrant that the **SHEAR** product will detect all hazardous shear conditions.
> When convective storms are in an air terminal area, there is danger of microburst. The **SHEAR** product is only one of many indicators that such a hazardous condition may exist. Normal precautions to avoid suspected wind shear should be used, even if the **SHEAR** algorithm does not detect shear.
> Vaisala shall not be liable for damages of any kind for failure of the **SHEAR** algorithm to detect hazardous wind shear or for false alarms that may occur from use of the SHEAR algorithm.

▶ 1. To create a new **SHEAR** product, select **Type > SHEAR**.

2. To load an existing product, select **File > Open**

3. In **Data:Display**, select `V Shear`.

4. In **EL Angle**, define the elevation angle of data you want to display.
   Wind shears associated with gust fronts and microbursts are usually low-level phenomena, so angles of 1° or less are typically used to detect these.

5. Define **Rng / Az Filter**.
   The first number defines the range length scale for the **SHEAR** product, in km. This distance determines the size of the skip that the **SHEAR** product can tolerate when differencing over missing data. It also specifies the length of the radial smoother, which is applied before the XY smoother.
   The second number is the azimuth length scale in degrees.
   The skipping and azimuth range length are analogous to the range filter.

6. In **Shear Type**, select which type of shear you want to compute:
   • Radial Shear
     The bin-to-bin difference in the radial velocity.
   • Azimuthal Shear
     The azimuth-to-azimuth difference in the radial velocity.
   • Elevation Shear
     The difference between the radial velocity at the selected and next higher elevation angle. Requires a volume scan.
   • Combined Shears
     The RMS value of the radial, azimuth, and/or elevation shear. This is sometimes referred to as the total shear.

7. In **VVP**, select the name of a **VVP** product.
   The wind speed and direction from the nearest **VVP** wind product of this name is used to remove the effects of the mean wind on the azimuth shear calculation.
   This **VVP** product must be scheduled to run. Use of the azimuth or combined shear is not recommended without this correction.

8. In **VVP Age**, specify the maximum number of minutes between the data for the **VVP** product and the data for this product.
   This prevents old **VVP** winds from being misapplied. If a **VVP** product is not found in this time window, the mean wind correction is not applied when computing the azimuth shear.

More Information

‣ RHI: Range Height Indicator (page 68)
‣ Display Parameters (page 29)
‣ Product Parameters (page 25)
‣ Associating Products with Tasks (page 19)

5.6.1 SHEAR Algorithm

The **SHEAR** algorithm shown in the following figure is for radial shear. The azimuth shear is similar in "B-scan" space (range-azimuth).

Input Radial Velocity Field

| V_1 | V_2 | | V_4 | | | $V{\sim}0$ | $V{\sim}0$ | V_9 | V_{10} | | V_{11} |

Reject V ~ 0 Bins

| V_1 | V_2 | | V_4 | | | | | V_9 | V_{10} | | V_{11} |

Difference Adjacent Bins

Gap > 0.5 *Range Filter Length

| ΔV_1 | ΔV_2 | | | | | | ΔV_9 | ΔV_{10} | |

Filter and Smooth

1/2

1/4

$\longrightarrow$ Convolve Triangular Filter

| $\Delta V'_1$ | $\Delta V'_2$ | | | | | | $\Delta V'_9$ | $\Delta V'_{10}$ | |

$$\Delta V'_2 \;=\; \frac{\frac{1}{4}\Delta V_1 \;+\; \frac{1}{2}\Delta V_2 \;+\; \frac{1}{4}0}{W_2}$$

$$where \quad W_2 \;=\; \frac{1}{4}\delta_1 \;+\; \frac{1}{2}\delta_2 \;+\; \frac{1}{4}\delta_3$$

and $\quad \delta_n = 0 \quad$ if $\quad \Delta V_n$ is thresholded, else $\delta_n = 1$

If the weight W is $\;\leq\; \frac{1}{2}$, then bin is thresholded (eliminates speckle bins).
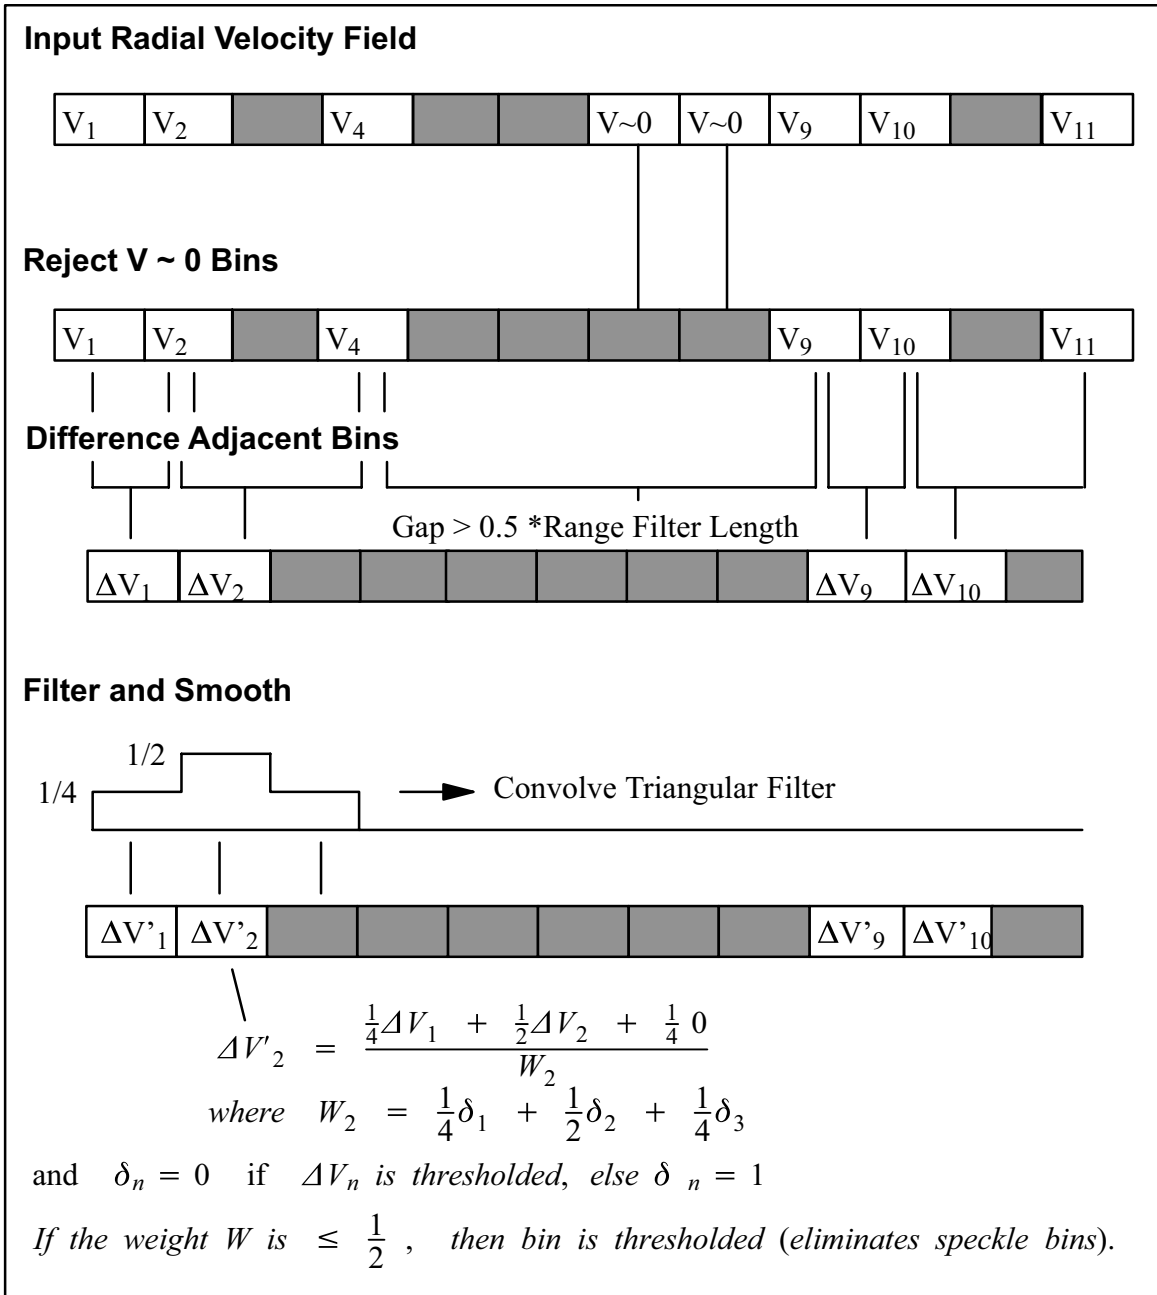
Figure 63   Radial Shear Algorithm Schematic

Point Clutter Bin Removal

Point clutter bins produce erroneous high shear values when they are differenced with surrounding valid velocities. The first step in the shear algorithm is to remove obvious clutter bins which were not canceled by the signal processor clutter filter. All velocity bins having a velocity within 2 % of 0 velocity (as compared to the full velocity range) are tagged as potential Clutter Bins. A Clutter Length Scale is defined as 1/3 of the Range Filter Length value in km. Contiguous runs of Clutter Bins that are less than or equal to the Clutter Length Scale are rejected.

For example, suppose the bin spacing and Range Filter Length values are selected so that Clutter Length Scale corresponds to 3 range bins. Then any single isolated Clutter Bin would be rejected, as would runs of 2 or 3 clutter bins. Runs of 4 or more clutter bins would not be rejected because they are interpreted as real weather.

Regardless of the Clutter Length Scale, isolated clutter bins (single bins) are always removed.

Differencing (radial difference example)

The next step is to perform the range differencing of the radial velocities. The radial wind shear algorithm computes the shear by taking the bin-to-bin difference on a **PPI** surface.

If the differencing algorithm encounters a blank bin, it skips out in range to the next valid bin to take a difference, provided the bin is within 0.5*Range Filter Length.

For example, if the Range Filter is set to correspond to 3 bins, the differencing algorithm differences over a gap of one bin, but not over a gap of 2 or more bins. The velocity differences are placed as close as possible to the center of the difference interval, to the nearest bin that is less than or equal to half of the difference interval.

Smooth the Differences

Differencing is inherently noisy. The next step is to smooth the computed differences. This not only reduces the noise, it also fills in small gaps of missing bins and cancels any remaining isolated bins.

The smoothing algorithm uses a triangular window whose total width is equal to the Range Filter Length value. The sum of the weights defined by the triangular filter is always 1 unless a bin is thresholded.

In the previous figure, the Range Filter Length corresponds to 3 bins. Each velocity in the filter range is multiplied by the corresponding weight and then summed. The mean is obtained by dividing by the total weight. Thresholded bins are not weighted. If the total weight is less than or equal to 0.5, the shear value is rejected. This eliminates isolated bins. This approach can also fill gaps. A longer range filter rejects a longer run of isolated bins and fills a longer gap.

Azimuth Shear Calculation

The following figure is a depiction of "B-Scan" space (a **PPI** surface). The coordinates are range and azimuth. The radial shear is obtained by differencing in range along a radial.
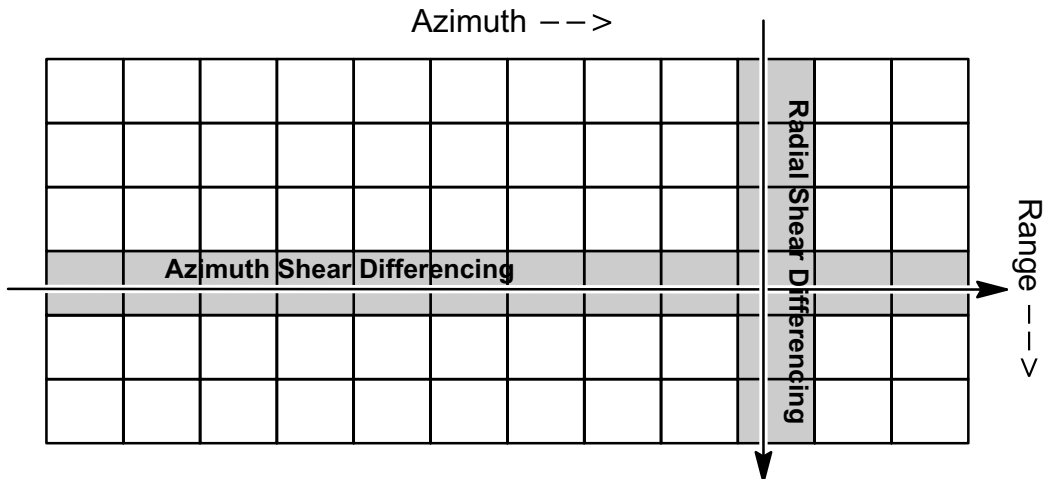
Figure 64   B-Scan Space (**PPI** surface)

The azimuth shear algorithm is analogous to the radial shear algorithm except that:

- Differencing is performed from ray-to-ray at constant range.
- The azimuth filter length is specified in degrees.
- The final azimuth values are range normalized and corrected for the apparent azimuth shear caused by the mean wind. The mean wind value is obtained from the **VVP** algorithm. This correction is an optional step.

For information on the mean velocity correction, see 5.6.2 Optimizing SHEAR for Microburst Detection (page 189).

Shear Magnitude

After both the radial and azimuth shears have been obtained, the next step is to compute the shear magnitude. This is the square root of the sum of the squares of the 2 shear values (radial and azimuth) associated with each bin.

If only a single shear is selected (radial or azimuth), this step is skipped.

Convert to Cartesian and Optional Final Cartesian Smoothing

The conversion to Cartesian transforms the B-scan values of the shear magnitude to X-Y coordinates.

The algorithm is identical to the **PPI** algorithm, which uses a last-in approach for both range and azimuth filling. An optional 2D smoother with selectable length scale can be passed over the data at this point.

5.6.2 Optimizing SHEAR for Microburst Detection

Accurate detection and timely reporting of microbursts in an air traffic control environment requires that the system be properly configured and optimized.

Vaisala recommends you work with a qualified consultant to assist with tuning and evaluating the IRIS **SHEAR** algorithm, data acquisition, radar siting and warning reports.

> The optimization steps presented here are only guidelines and should not be substituted for a thorough evaluation of your particular site. IRIS provides many of the tools needed to perform such an evaluation.

Table 18   Mean Velocity Shear Correction

The radial velocity can be expressed in terms of the u and v components of the wind (the vertical velocity is ignored here because the observations are typically near horizontal:
where: $V_R = u\cos\theta\sin\varphi + V\sin\theta\sin\varphi$ V_r is the radial velocity u is the x-component of the wind (East) v is the y-component of the wind (North) θ is the azimuth defined as positive CCW from the x-axis (East) φ is the elevation defined as positive upward Let the wind be represented by a mean wind [`u`_{`o`}`, v`_{`o`}] and a perturbation wind [`u'`, `v'`], then $V_R = \left(u_o + u'\right)\cos\theta\sin\varphi + \left(v_o + v'\right)\sin\theta\sin\varphi$ The azimuth derivative is then, $\frac{\alpha V_R}{\alpha\theta} + -u_o\sin\theta\sin\varphi + v_o\cos\theta\sin\varphi + \frac{\alpha}{\alpha\theta}[u'\sin\theta\sin\varphi + v'\cos\theta\sin\varphi]$ where the first term is the apparent azimuth shear caused by the mean wind. Thus, when the azimuth derivative is computed, the correction term is applied as follows: $\frac{\vartriangle V_R}{\vartriangle\theta} + u_o\sin\theta\sin\varphi - v_o\cos\theta\sin\varphi = \frac{\alpha}{\alpha\theta}[u'\cos\theta\sin\varphi + v'\sin\theta\sin\varphi]$ The values of u_o and v_o are obtained from the **VVP** algorithm. In terms of the azimuth `#beta#` measured in the traditional radar sense (clockwise from north) the correction is: $\theta = 90 - \beta$ $u_o\cos\beta\sin\varphi - v_o\sin\beta\sin\varphi$

The **SHEAR** product is first produced from tasks that are optimized for close-range, high-resolution measurement of radial velocity.

The **SHEAR** products are then fed to the **WARN** product, which checks to see if the strength and size of the wind shear regions exceed a threshold value in the protected areas.

Additional criteria, such as reflectivity aloft from a higher scan or a requirement to see the shear on two adjacent scans, can be added to the **WARN** product to reduce the false alarm rate if this is a problem for the particular location.

1. Setup an IRIS configuration for wind shear, so that when potentially hazardous weather approaches the terminal area, IRIS can be switched to wind shear monitoring mode by loading the configuration into the **Radar Status** menu.
   You can a operate the **WARN** product in surveillance mode to alert the operator that the mode should be switched, or the automatic mode switching feature can be used. For example, the surveillance mode could consist of a 15 elevation volume scan sequence. The **WARN** product can be keyed on a severe storm indicator such as **VIL**. If **VIL** exceeds a threshold value typical of severe storms for the area, and this **VIL** is within 30 km (18.6 mi) of the terminal, the **WARN** product can issue an alarm advising the operator or the automatic mode switch to change to the wind shear detection mode.

2. Optimize the **TASK Configuration** menu for wind shear detection as shown in the example in the following figure.
    This example uses high-resolution sampling in range (125 or 62.5 m (205 ft 1 in)). Microbursts are low-level phenomena, so it is best to configure elevation angles that are as low as can be tolerated by the surrounding clutter.
    Typically an angle set such as 0.5°, 1.0° and 5.0° is useful. The bottom 2 angles can be used for shear detection, while the upper angle can serve as the basis for an additional reflectivity aloft criterion for the **WARN** product. Sector scan mode is used to limit the data coverage to the terminal area which speeds the update rate.
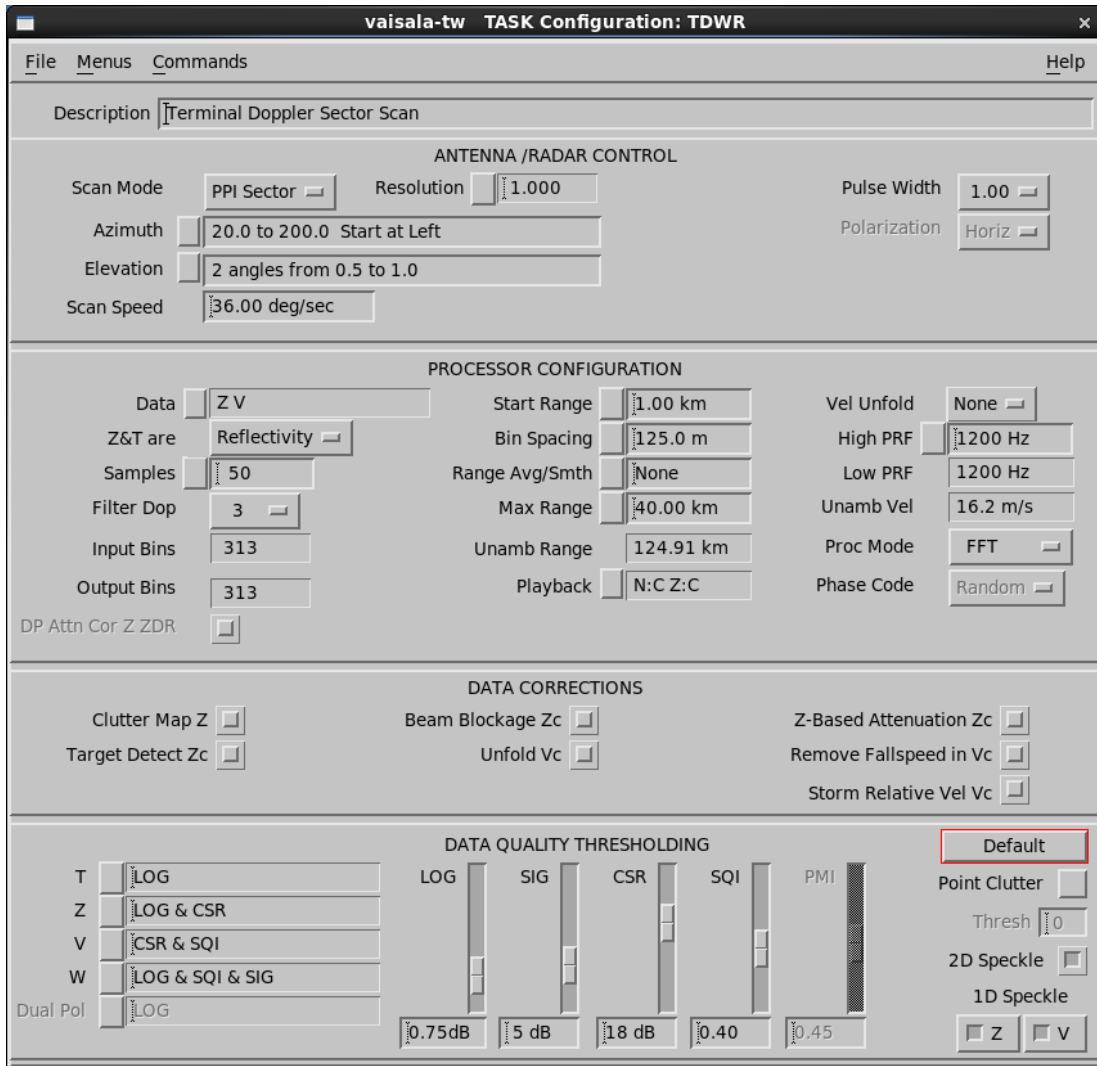
Figure 65   TASK Configuration for Microburst Detection

3. Set the **TASK Scheduler** in wind shear detection mode to provide an update rate (repeat time) of at least one complete task per minute.
    To provide optimal response, do not schedule other tasks. This allows IRIS to devote its full resources to detecting microbursts.

4. Optimize the **SHEAR** product configuration:
   - Set the **Product Range** field to the maximum range of interest, and should roughly correspond to the input bin spacing.
     The 30 km (18.6 mi) range and 480 by 480 pixel product resolution in the example match the 125 m (410 ft 1 in) sampling in the task.
     In the example elevation angles of 0.5, 1.0, and 5.0, separate shear products can be configured for each of the lowest 2 angles.
   - It is recommended not to use any additional XY smoothing when the **SHEAR** product is used for microburst detection because smoothing tends to diminish the peak values that are of interest in the shear calculation.
     The radial smoothing performed by the range filter, and the subsequent smoothing performed in the **WARN** product are adequate.

   The **WARN** product looks at the **SHEAR** products to determine whether to issue a warning. In the example of 2 low-level scans (0.5º and 0.9º), each with its own **SHEAR** product, the **WARN** product can be set to examine each. A threshold level of 8 m/s/km is a good indicator of hazardous wind shear. A 1 km (0.6 mi) area threshold can be used which corresponds to roughly 64 pixels for the case of 125 m (410 ft 1 in) bin spacing, 30 km (18.6 mi) range and 480 × 480 product resolution.

5. Optimize the **Product Output** menu to send the **WARN** product to a workstation so that personnel can view the situation display.

   a. Make the automatic output request so that the shear product is sent every time that it is made.

   b. Make sure a workstation is running the IRIS menus so that when windshear is detected, the audible warning message and text can be viewed.

   Warning messages are issued almost instantaneously when wind shear is detected. Within a few seconds of receiving a warning message, the situation display is presented. This depends on the speed of the communications link connecting your display. However, because the situation display is an overlay product without a complex picture, it can be transmitted efficiently.

6. Tune the wind shear detection algorithms for each site.
   This requires that qualified personnel perform the necessary optimization for your particular meteorological and operational environment.
   You must go through a thorough evaluation phase to test your system on your particular site for your particular weather.
   It is recommended that you make and record a **RAW** product every time a microburst is detected and perform an analysis to determine if it was a real event or a false alarm.
   To obtain a larger sample, you may want to expand the protected area to include all the total area around the radar to the maximum range of detection required (typically 20 to 30 km (12.4 to 18.6 mi)) for air terminal applications. A qualified meteorologist should assist with this evaluation.

7. If it is determined that false alarms are occurring, there are several approaches you can take to balance sensitivity against false alarm rate.
In tuning your system to reduce false alarms, it is important to make gradual changes and evaluate them on test cases that you have recorded. Some examples of actions that you can take to reduce false alarms are as follows (note the reverse of each example can be used to increase sensitivity):

- Gradually increase the range filter in the **SHEAR Product Configuration** menu from 1 to 2 km (0.6 to 1.2 mi).
- Increase the **SHEAR** magnitude threshold in the **WARN Product Configuration** menu.
- Increase the threshold AREA size in the **WARN Product Configuration** menu up to 2 km (1.2 mi).
- Require that 2 low-level scans from the same task both detect the **SHEAR** in the **WARN** product: a two-look test.
- Add a warning criterion that requires reflectivity aloft. This can be based on a dBZ PPI at a higher angle.
  Make sure dry microbursts are not common in your area. If they are, you may need to use a `VIL` criterion and employ more elevation angles in your associated task.

5.7 **SLINE**: Shear Line

Figure 66   **SLINE** Example Configuration

SLINE (shear line, or frontal boundary) marks the transition between two air masses on the image.

Because the wind can be different on either side of the air mass, sudden wind shifts occur when a frontal boundary passes.

The following figure shows an example of a typical shear line. The most severe wind shifts are associated with mesoscale gust fronts, which occur at the leading edge of a thunderstorm outflow, and with larger scale cold fronts, which may be associated with severe weather.

Figure 67   Typical Shear Line Example

The **SLINE** product can detect sudden wind shifts and provide advance warning of their approach, with best estimates of the arrival time and the expected speed and direction of the new wind.

1. Select **Type > SLINE**.

2. To load an existing product, select **File > Open**

3. In **Data:Display**, select `V Shear`.

4. In **EL Angle**, define the elevation angle of data you want to display.
   Wind shears associated with gust fronts and microbursts are usually low-level phenomena, so angles of 1° or less are typically used to detect these.
   You may select 1 or 2 angles. If you select 2 angles, make sure both are above the threshold to be included in the shear line.

5. Define **Rng / Az Filter**.
   The first number defines the range length scale for the **SHEAR** product, in km. This distance determines the size of the skip that the **SHEAR** product can tolerate when differencing over missing data. It also specifies the length of the radial smoother, which is applied before the XY smoother.
   The second number is the azimuth length scale in degrees.
   The skipping and azimuth range length are analogous to the range filter.

6. In **Shear Type**, select which type of shear you want to compute:
   - Radial Shear
     The bin-to-bin difference in the radial velocity.
   - Azimuthal Shear
     The azimuth-to-azimuth difference in the radial velocity.
   - Elevation Shear
     The difference between the radial velocity at the selected and next higher elevation angle. Requires a volume scan.
   - Combined Shears
     The RMS value of the radial, azimuth, and/or elevation shear. This is sometimes referred to as the total shear.

7. In **VVP**, select the name of a **VVP** product.
   The wind speed and direction from the nearest **VVP** wind product of this name is used to remove the effects of the mean wind on the azimuth shear calculation.
   This **VVP** product must be scheduled to run. Use of the azimuth or combined shear is not recommended without this correction.

8. In **VVP Age**, specify the maximum number of minutes between the data for the **VVP** product and the data for this product.
   This prevents old **VVP** winds from being misapplied. If a **VVP** product is not found in this time window, the mean wind correction is not applied when computing the azimuth shear.

9. Define **SLINE** settings:

   Count and Time

       Enter the number of forecast points and the time interval between each forecast point and select **Ok**.
       The time you specify is inserted in the field.

   Max Time Step

       The forecast of shear line motion uses the previous shear line product.
       Max Time Step specifies how far back in time the previous **SLINE** product can be. If none exists in this time frame, no forecast is made.

   Max Velocity

       Enter the maximum expected velocity. If the computed velocity exceeds this value, the shear line is considered to be a new feature, and no forecast is made.

   Shear Threshold

       Enter the minimum speed difference per km that must be present to indicate a shear line. Winds that do not meet this speed are ignored.

   Threshold Area

       Enter the minimum size of the wind shear area in square km that must be present to indicate a shear line. Wind shear areas below this threshold are ignored.

   RMS Fit Tolerance

       Enter the standard deviation allowed when fitting the shear line. The line is discarded if the points do not fit within this tolerance, indicating a random scattering of points and not a shear line.

   Minimum Sline Length

       Enter the minimum length distance in km used to estimate a shear line. Shear lines that are below this length are ignored.

10. Configure **SLINE** for use with warnings.

   Protected Areas for Warning Alert

   Specify one or more protected areas. A warning is generated if a shear line is within a protected area or is forecast to move into one.

   Select **Protected Areas** to see a list of areas and toggle a choice on or off. When you exit the list, the names of the selected areas are displayed in the field.

   Protected areas are configured with the **Setup** utility. They are constrained to be rectangles with arbitrary size and orientation angle. Up to 32 areas can be configured.

   If you make a change to **Setup**, you must restart IRIS for the change to take effect. See *IRIS and RDA Utilities Guide*.

   TDWR Style

   IRIS supports two formats of warning messages.

   In TDWR format, only the strongest centroid in the highest priority area is reported along with its strength. For example: `MBA 3MF 30K-`, in spoken language: `Microburst Alert, Three Mile Final, 30 Knot loss`.

   The older IRIS format reports all centroids in all protected areas, for example: `3 MBA warnings at 11:30 in: 3MF, 2MF`.

   The messages are displayed at the bottom of the display screen, and optionally spoken.

   Say Warning

   When enabled, the product generator to audibly signals the warning message in addition to displaying it as a message.

   You must also select **Setup > Enable Speech** in IRIS.

11. To troubleshoot or test **SLINE** products, enable a thresholded version of the input **SLINE** product by selecting **Make Diagnostic Results**.

   This has the same product name as the **SLINE**.

More Information

‣ Display Parameters (page 29)
‣ Associating Products with Tasks (page 19)
‣ Product Parameters (page 25)

5.7.1 Shear Line Radar Signatures

Detection of shear lines is based on the Doppler velocity rather than the reflectivity. However, accurate measurement of the Doppler velocity requires sufficient returned power. The radar reflectivity associated with shear lines can be from precipitation or clear air echoes. Because clear air echoes tend to be weak, the most efficient detection requires that the radar scan slowly for sufficient averaging.

The climatology of shear lines for a particular location dictates the scan strategy. If shear lines tend to be associated with weak returns, a slow scan rate is necessary for reliable detection. The trade-off between reliable detection of weak echo features and the need for rapid updates to detect phenomena such as microbursts is a decision that you must make, based on your site and application.

Shear lines are characterized by both vorticity and convergence (negative divergence). The vorticity associated with cold fronts in the Northern hemisphere is typically positive. The vorticity associated with thunderstorm gust fronts is usually, but not always, positive. Shear lines are not characterized by positive divergence, as are microbursts.

Both azimuthal shear (partial vorticity) and radial shear (partial divergence) typically show a signature for a shear line.

5.7.2 Shear Line Algorithm

The Shear Line (**SLINE**) algorithm uses the uniform wind assumption.

1. The radial, azimuthal, or combined shear are computed using the **SHEAR** algorithm.
   The previous **VVP** wind is used to remove the apparent azimuth shear caused by the mean wind.
   The shears are computed using polar data and the output is then X-Y data to a selectable grid spacing.
2. The shear magnitude is thresholded with regard to strength and area within a specified area.
   Shear regions less than the threshold or smaller than the area limit are not considered.
3. The shear magnitude values of 2 elevations are ANDed.
   This step is optional if a second elevation is specified.
4. A curve is fitted to the shear field to detect a linear region.
   Points that fall outside the curve are eliminated, and the curve is refit.
   Only one shear line can be considered in a product.
5. The speed of motion normal to the gust front is obtained by comparing the data from the current shear line location with the data from the previous location.
6. The **VVP** wind is computed in both areas defined by the shear line using the uniform wind assumption.
7. The computed motion is projected to determine if and when the shear line will cross each protected area.
8. The results are output.

5.8 **SWS**: Surface Weather Stations

IRIS can store and display METAR and SYNOP reports from Surface Weather Stations. These are stored in the **SWS** product.

A single product file contains reports from many (or all) stations in the radar coverage area. Since this data is not generated from the radar data, there is no product configuration menu to create such products. Instead the product file is created when METAR or SYNOP text messages from a surface weather station are passed through an IRIS pipe program.

This data is displayed as a series of "station plots", which show various stored data, such as the temperature, wind direction, pressure, and similar. These can be plotted stand-alone, or on top of radar products. You can also click the cursor to see the raw station reports.

5.8.1 Displaying SWS Data
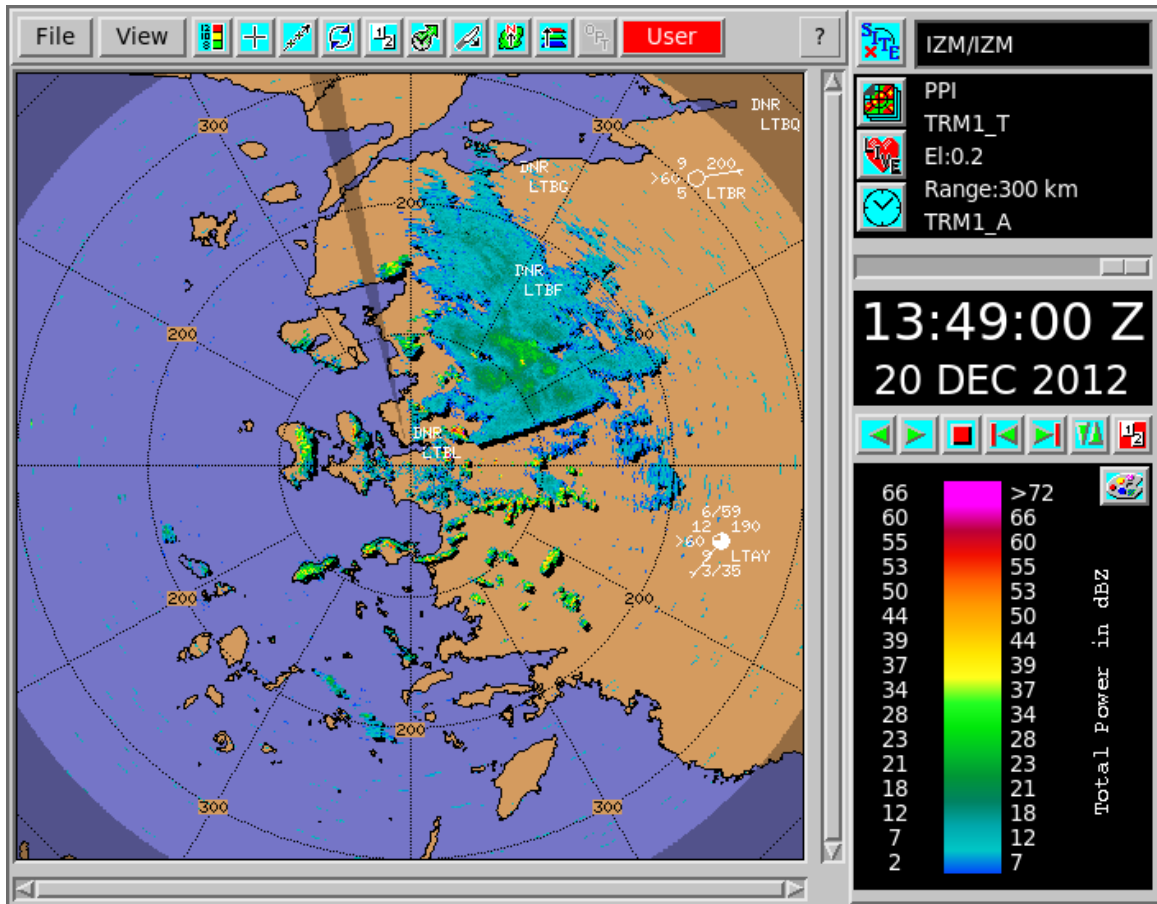
Figure 68   Example of Radar Data with an Overlay of SWS Station Plots

In the previous example, the individual station plots are drawn centered at the station locations. All the stations in the display coverage area are shown.

In the IRIS station plot format, the plot is standardized, although different reporting stations vary in measuring capabilities as well as in data the span and resolution. Individual numbers or graphics are not shown on the plot if a station does report or if the data is turned off in the output options.

Figure 69   IRIS Single Station Plot

The plot details are as follows.

#1 Cloud Coverage

A pie chart showing coverage in oktas (that is, eights) of sky coverage drawn using standard symbols. The example shows 6 oktas.

Because METAR messages do not give cloud cover in oktas units, METAR text is mapped as follows:

METAR message	Oktas
SKC or **CLR**	0 oktas
FEW	1 okta
SCT	3 oktas
BKN	6 oktas
OVC	8 oktas

Symbol	Description	Symbol	Description
◯	Sky clear (0 oktas)	◖	6 oktas of sky covered
◐	1 okta or less of sky covered, but not zero	◑	7 oktas of sky covered
◔	2 oktas of sky covered	●	8 oktas of sky covered
◑	3 oktas of sky covered	⊗	Sky obscured by fog or other meteorological phenomena
◑	4 oktas of sky covered	⊖	Cloud cover obscured for other reasons or not observed
◑	5 oktas of sky covered		

Figure 70   Okta Standard Symbols

#2 Wind Barb

A line with feathers showing the wind speed and direction, to the nearest 5 knots.

A knot (nautical mile per hour) is close to half a meter/second, so 20 knots would be 10 m/s. The line points up wind. It has a half feather line for 5 knots, and full length lines for 10 knots, and a triangular flag for 50 knots. If the wind speed is 0, then a circle is draw outside the sky coverage circle.

The example shows a wind of 5 knots from the south.

#3 Present Weather

There is a large set of available symbols, shown in the following figure. The example in Figure 68 (page 200) shows 2 dots, which means light rain.

Matching of METAR present weather text to symbols in Table 19 (page 203) is not necessarily endorsed by the National Weather Service (NWS) or the World Meteorological Organization (WMO).

In the following figure, blue numbers in upper-left corner of white boxes indicate the priority for plotting in event more than one symbol is possible (symbols in gray boxes have no corresponding METAR present weather text). For graphical representation of METARs using this table, see http://www.rap.ucar.edu/weather/surface/

Figure 71   Present Weather (METAR text-to-symbol matching)

#4 Visibility

This code is a 2-digit shorthand for a somewhat logarithmic scale, as shown in Table 19 (page 203).

The ">60" in the example means greater than 10 km (6.2 mi).

METAR reports have a maximum value of "greater than 10 km", so you there are many ">60" on the displays.

Table 19   METAR Visibility Codes

Code Value	Meaning
00	Less than 100 meters (328 ft)
01 ... 50	Code * 100 meters (100 to 5000 meters (328 to 16 404 ft))
51 ... 55	Unused
56 ... 80	(Code-50) * 1000 meters (3280 ft) (6 to 30 km (3.7 to 18.6 mi))
81 ... 88	(Code-74) * 5000 meters (16 404 ft) (25 to 70 km (15.5 to 43.5 mi))
89	>70 km (43.5 mi)
90 ... 99	Unused for land stations

#5 Temperature

Temperature (°C).

#6 Cloud Height

There can be up to 3 levels of clouds plotted in the format O/HH.

- The lowest level clouds are drawn below the station, in this case "3/40"
- The second level of clouds is drawn above the station, which is "6/60" in the example
- There can be a third level drawn above the second, not shown in this example.

The first number is the oktas of coverage, and the second number is a code value for the height. These code values are similar to the visibility code values, except the units are feet, as shown in the following table.

The example shows 3/8 coverage of clouds at 1219 m (4000 feet), and 6/ 8 coverage at 3048 m (10000 feet). SYNOP reports contain only 1 level of cloud height, so the others are not seen.

Table 20   Cloud Height Codes for **SWS**

Code Value	Meaning
00	Less than 100 feet
01 ... 50	Code * 100 feet (100 to 5000 feet)
51 ... 55	Unused
56 ... 75	(Code-50) * 1000 feet (6000 to 25000 feet)
76 ... 99	Unused

#7 Sea Level Pressure

Sea Level Pressure in 1/10 of a millibar, with the leading "9" or "10" removed.

If the string starts with a 0 through 4, prefix a "10", otherwise prefix a "9". In the example, this is 1015.0 millibars.

METAR reports the pressure only to the nearest millibar, so all the numbers end with "0". This pressure is the QNH pressure, which is corrected to sea level.

#8 Pressure Trend

Shows the unsigned change in pressure over the last 3 hours in 1/10 of a millibar. It is followed by a graphic showing the change trend.

For example, "4/" means that the pressure went up by 0.4 millibars in the last 3 hours.

METAR reports do not contain pressure trend, so this is blank.

The following figure shows the available symbols.

Symbol	Description of curve	Pressure now compared with 3 hours ago
∧	Rising, then Falling	Higher
⌐	Rising, then steady	Higher
/	Rising	Higher
√	Falling, then Rising	Higher
\	Falling, then Rising	Lower
L	Falling, then steady	Lower
\	Falling	Lower
∧	Rising, then Falling	Lower

Figure 72   **SWS** Pressure Trend Symbols

#9 Station ID

- METAR stations use a 4-character string.
- SYNOP stations use a 5-digit number.

#10 Dew Point

Dew Point (°C).

5.8.2 SWS Output Options

When displaying an **SWS** product in the **Quick Look Window**, you can configure what data is displayed using the **SWS Tool**.

You can start **SWS Tool** from **Quick Look Window > OPT** or from **Product Output Menu > Opts**.

For all selected data fields, IRIS displays the data and stores the data in the file.

> **Pressure** includes both the current pressure and the pressure trend.
> **Ceiling** includes (up to) 3 cloud level heights.

Figure 73   **SWS Tool**

5.8.3 SWS Raw Report

When displaying an SWS product in the **Quick Look Window**, you can hover over a station and right-click the mouse to show a report of the raw METAR or SYNOP reports for that station.

If there are more than one report in the product file, it displays all, one on each line. The station plot shows only the most recent report in the product.
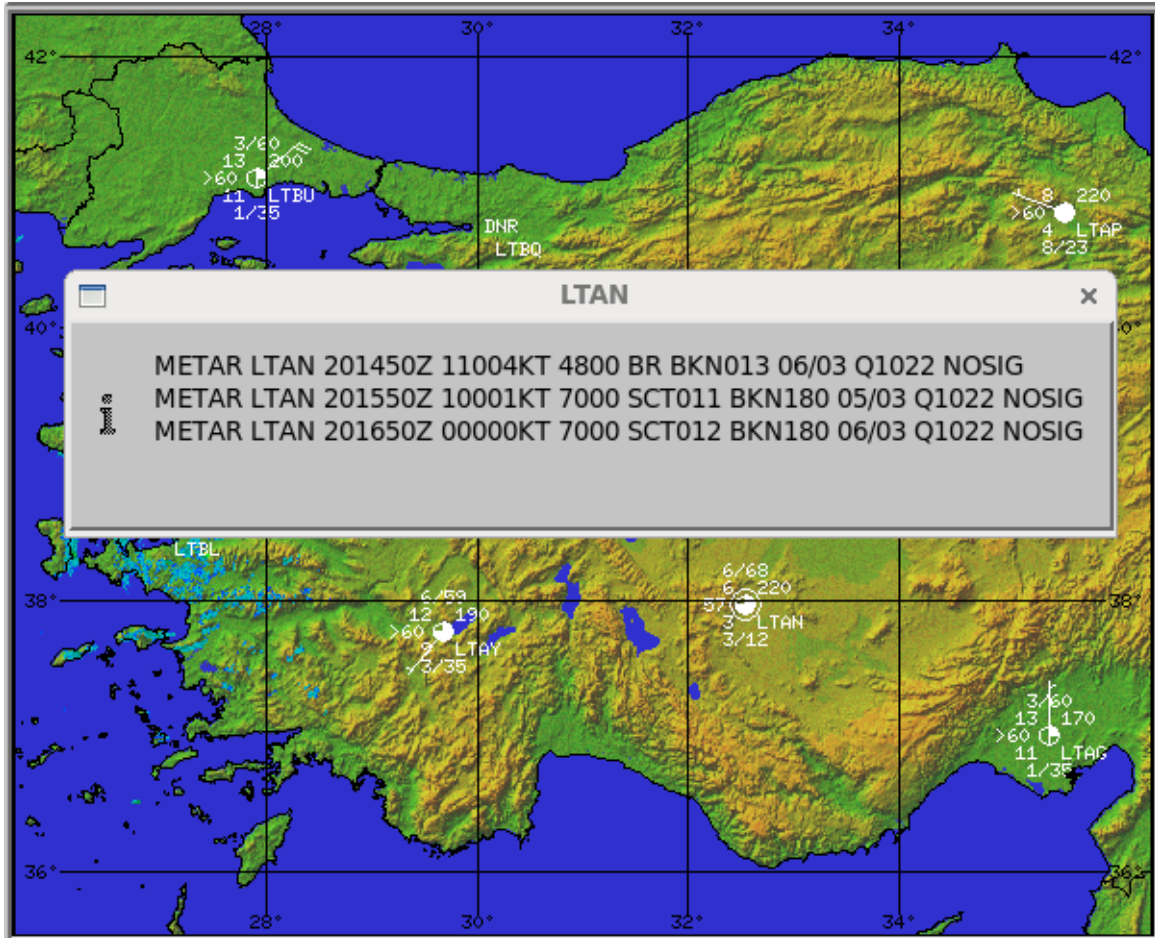
Figure 74   **SWS** Raw Report Example

5.8.4 Configuring SWS Input

SWS data arrives on an IRIS system through a file transfer and ingest mechanism.

You must configure a source of METAR and/or SYNOP messages to transfer ASCII files to a known directory on the IRIS Analysis system. IRIS monitors this directory for new files, then runs an input pipe to convert this data to IRIS's internal **SWS** format and insert into the IRIS product directory.

You must configure the **SWS** input.

▶ 1. Make sure you use a recent version of the `metar2iris` pipe by copying the binary and config files from *config_template* directory to *config* directory.
This pipe parses metar and synop messages.

```
$ cd /usr/sigmet/config_template
$ cp pipes/metar2iris /usr/sigmet/config/pipes|
$ cp metar* /usr/sigmet/config
```

2. Configure and save a new input device in **Setup > Input**. For example:
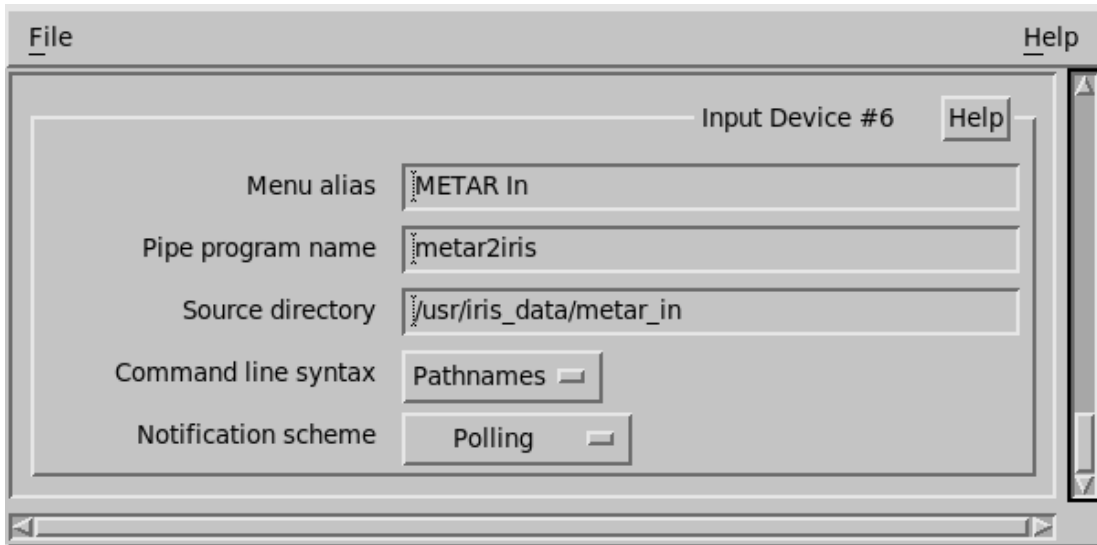
3. Edit the following configuration files:

File Name	Description

`metar_site.conf`
 	
`synop_site.conf`
	Configures the METAR or SYNOP station site list consisting of ICAO site ID, (optional WMO site ID), LAT/LON, altitude, description, and priority.
	
	Since the METAR/SYNOP texts do not include the latitude and longitude of the stations, this file is essential.
	
	When a METAR or SYNOP file arrives, all stations in this file are expected.
	
	• Any station listed in the file, but missing from the report, are marked as **DNR** (Did Not Report).
	• Any station in the report, but missing from the file is deleted, and a message is placed in the pipe log file.
	
	You must maintain the files to remove old stations, and add new ones.
	
	For a rough list of METAR and SYNOP stations, see: http://weather.rap.ucar.edu/surface/stations.txt

`fmetar2Iris.conf`	Configures the product's geographical bounding box. This is the size of the box when the **SWS** product is displayed alone. Any station outside this box is discarded. We also configure the **SWS** product name, and IRIS site name, and similar.

The template files are configured for Turkey.

4. Create the input directory. These directories must be owned by operator.

```
$ su
# mkdir /usr/iris_data/metar_in
# chown operator:users /usr/iris_data/metar_in
```

5. Restart IRIS.

5.8.5 Feeding In SWS Data

▶ 1. For testing, manually feed in example data to the */usr/iris_data/metar_in* directory.

2. For testing, download current METARs from: http://aviationweather.gov/adds/metars/. You must edit the file to add the GTS header, and METAR on each line.

3. Check that you get an **SWS** product in IRIS.

4. Check for warnings in the pipe log file *metar2iris.log* in */usr/iris_data/log*. The time stamps consist of the day of the month, followed by hour and minute. The pipe must guess the month and year, based on the current time.
IRIS expects to receive one ASCII file containing all SYNOP or METAR messages from the past 1 hour time. METAR special reports can also be sent into IRIS in-between hours.

5.8.6 METAR and WMO File Format Examples

METAR ASCII file format, including a WMO GTS header line and then each station's report. Each report has a separate time stamp. "=" signs are optional.

```
SATU20 LTAA 160150

METAR LTAB 160150Z VRB01KT CAVOK 14/06 Q1015 NOSIG RMK RWY06 VRB02KT=
METAR LTAD 160150Z 36003KT CAVOK 15/07 Q1015 NOSIG RMK RWY29 VRB02KT=
METAR LTAE 160150Z 35003KT CAVOK 12/04 Q1015 NOSIG RMK RWY21 02006KT=
METAR LTAG 160150Z 10003KT 9999 SCT035 25/22 Q1005 NOSIG=
```

SYNOP ASCII file format, including a WMO GTS header line, an **AAXX** line, then each station's report. All reports share the same time stamp.

```
SMTT60 LTAA 160000

AAXX 16004
17015 3//// /1605 10171 20168 30109 40120 57006 60001
17018 3//// /3101 10160 20069 38899 40112 57009 60001
17020 3//// /0402 10153 20147 30085 40124 56003 60001
17022 3//// /1203 10197 20120 39966 40120 57008 60001
```

6. Quick Look Window

6.1 Overview of Quick Look Window

The Quick Look Window (QLW) displays IRIS products automatically or by request so you can:

- Select product images for single-frame browsing, looping, and slide show.
- Select range rings, overlays, shadowing, and color scales.
- Use the cursor for position read-outs in radar and latitude/ longitude coordinates.
- Track, annotate, and forecast radar echo features, such as typhoons or severe convective storms.
- Shift radar displays forward in time.

Live IRIS allows you to interactively generate and display many radar product types in the QLW without performing the product configuration, scheduling, and output steps.
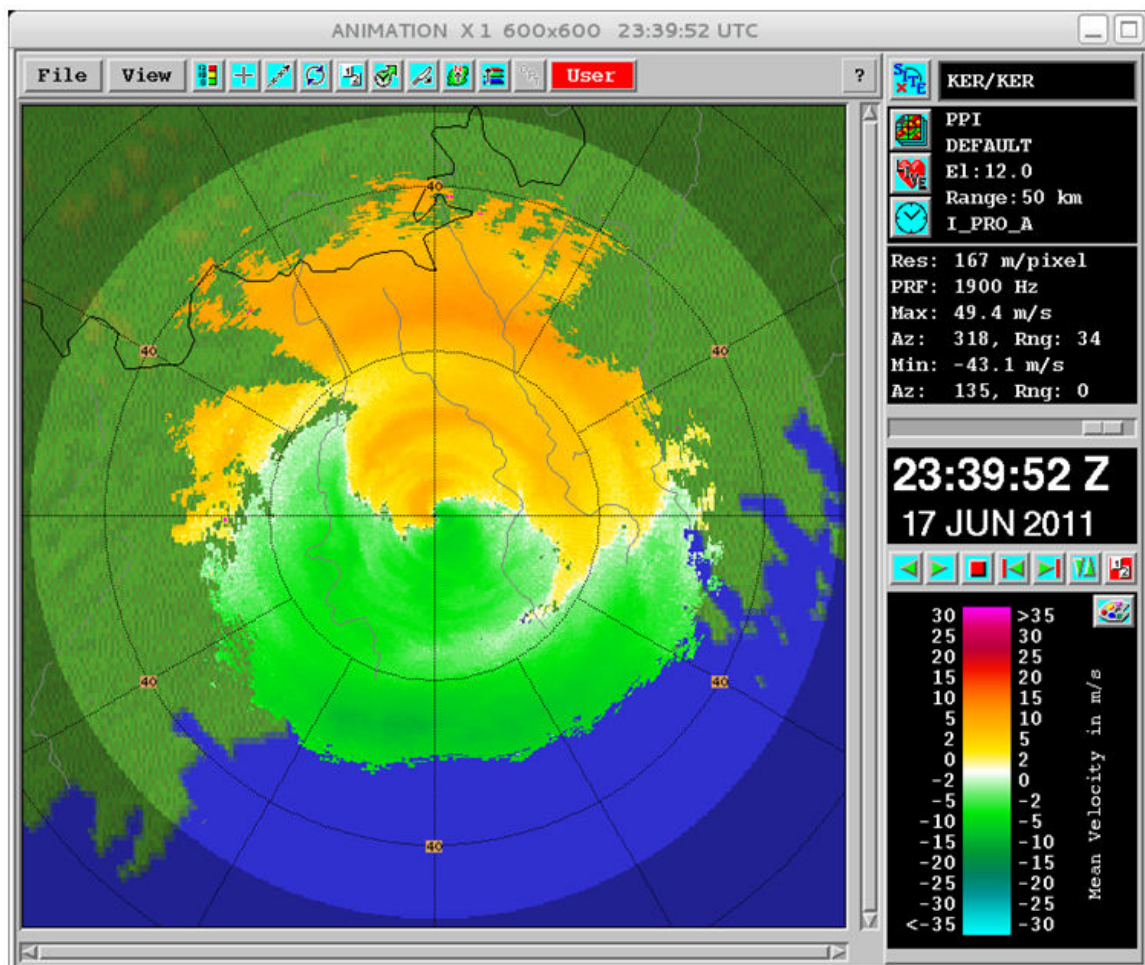
Figure 75   Example QLW: PPI Reflectivity (480 x 480 window size with range rings and overlay)

Title Bar

The title bar shows basic information, including:

- Name of the window that was assigned by the operator in **Setup > Output**.
  The example shows the name `wind #1`.
- Zoom level selected in the **View** tool.
  The example shows `X 1`, indicating no zoom.
- Window size in pixels selected in the **View** tool.
  This refers to the size of the image area. The actual window size is larger because of the borders and legends.
- The data time of the displayed image.

> When naming windows, consider including the product name. For example, if a **VIL** product is assigned by the **Product Output** menu to be displayed automatically in the window, name the window **VIL**.
> Also, consider naming the window to match the name of the workstation "work space" in the Common Desktop Environment (CDE), if this is supported on your workstation.

Tools

Icons below the title bar provide access to interactive functions such as loop configuration, track, cursor, forecast, and cross-section.

Select an icon to open the tool.

On the right of the tool area, the **QLW Mode** and the **QLW Indicator** alert operators to changes in the status of automatically assigned images in the QLW.

Legend

The legend provides tools for selecting radar sites, products and times, and displaying this information. You can also loop animations or browse single frames.

The color legend is displayed and labeled. The **Color Scale** allows you to select default color scales assigned by the operator, or to make your own color scales.

6.2 Setting-Up and Starting Quick Look Windows

You do not have to do anything to start the Quick Look Windows.

When IRIS starts, the windows that have been configured by your system manager (in **Setup > Output**) automatically appear on the screen. Most systems are set to have at least 2 windows, but you can have more (check with your system manager).

Windows can also be exported over the network from an IRIS workstation to other workstations that are running X-Windows, even computers running X-Window software on a Microsoft Windows OS. These workstations do not need to run IRIS (free displays). However, if you export displays over the network, your loop performance may not be as good as when you run loops on your local IRIS workstation. Your system manager can configure these remote windows.

> **i** If you log on to a system where IRIS is already running but there is no X session running (for example, screen lock or log out), the windows appear after a few minutes. It takes IRIS a little while to realize you are back.

After a window is up, there is no way to get rid of it except to minimize it where it is handy for quick viewing because weather images are always available and users cannot shutdown windows inadvertently. The minimize window shows a small thumbnail picture of the image in the display.

> **i** If a window crashes, radar operation and other IRIS functions continue normally and the window restarts automatically after a minute or two.

6.3 QLW Control and Monitoring

Quick Look Window Update Indicator

The Quick Look Window (QLW) update indicator alerts you to the status of the images displayed in the QLW. This allows you to see, even from a distance, that a window is set to display live data. The update indicator is next to the tool area **OPT** button.

Window Modes

The **Quick Look Window** operates in the following modes.

 (Green Background means automatic updating from IRIS)

- **Auto Mode**
  The image is updated automatically by the **Product Output Menu** (as set by the operator) with new images. This is non-interactive use of the window.

 (Red Background means no automatic updating from IRIS)

- **User Mode**
  In this mode, you decide what images are displayed in the window and can be actively manipulated (for example, creating a cross-sections or tracks without an interruption from new products arriving from the **Product Output** menu). Images are not automatically displayed from the **Product Output Menu**.

Slide (Green Background)

- **Slide Mode**
  Images are displayed based on the **Slide Show** tool configuration.

Loop (Green Background)

- **Loop Mode**
  Images are displayed based on the Loop/Animation tool configuration and the **Most Recent** button is checked so new images are being displayed as they arrive from IRIS.

Loop (Red Background)

- Images are displayed based on the Loop/Animation tool configuration and the **Most Recent** button is NOT checked so new images are not being displayed.

> You can select **User** and **Auto** modes through **Mode** in the tools area.
> **Loop** and **Slide** mode are selected automatically when a slide show or animation/loop starts.

In practice, you rarely need to manually switch the display mode, since IRIS does it automatically for you:

- When you first start to do something in the display, IRIS automatically switches from **Auto** to **User**. IRIS does not interrupt you by sending images while you are using the window.
- The QLW automatically switches to **Auto** mode after 5 minutes of inactivity.

Legend ON/OFF Button

Legend ON/OFF Button

The legend area, title bar and scroll bars may by toggled off by the **Legend** icon which is located on the left side of the tool section. Right-click the mouse in the display window to restore the legend or if you are using a touch screen double tap the upper left corner (100 × 100 pixels) of the window.

Turning off the legend if you want to reduce the size of the windows so that you can fit more windows on your screen.

The start-up location of each window and whether the legend is on/off is controlled in **Setup > Output**.

> With the legend turned off, you may want to enable the internal legend that shows a legend right in the display window instead of as a side menu.
> See 6.8 Display Options Tool - IRIS Product Output (page 220).

> Turning off the legend is useful for "tiling" the display windows such as for creating a multi-panel slide show.

Setting Operator Password

Some functions in the **Quick Look Window** are for operators only , since they can affect other Windows:

- Creating or deleting "home" reference points in the **Cursor** and **Track** tools.
- Creating or deleting names for virtual overlays (special combinations of various overlay layers) in the **Display Options** tool.

If you require access to these features, use the **File** tool to set the operator window password configured in **Setup > General**.

Select **Operator** and enter the password. Note that in general, the window password configured in **Setup > General** is not the same as the UNIX operator password.

> When you are done, select **File** and set the privilege back to **Observer**.

6.4 Viewing Products in Quick Look Window

In your workstation, you can view products in the **Quick Look Window**.

You can configure many display windows an IRIS system with the **Setup** utility.

You can send products to the display from either the **Product Output Menu** according to a product output schedule or from the **Quick Look Window**.
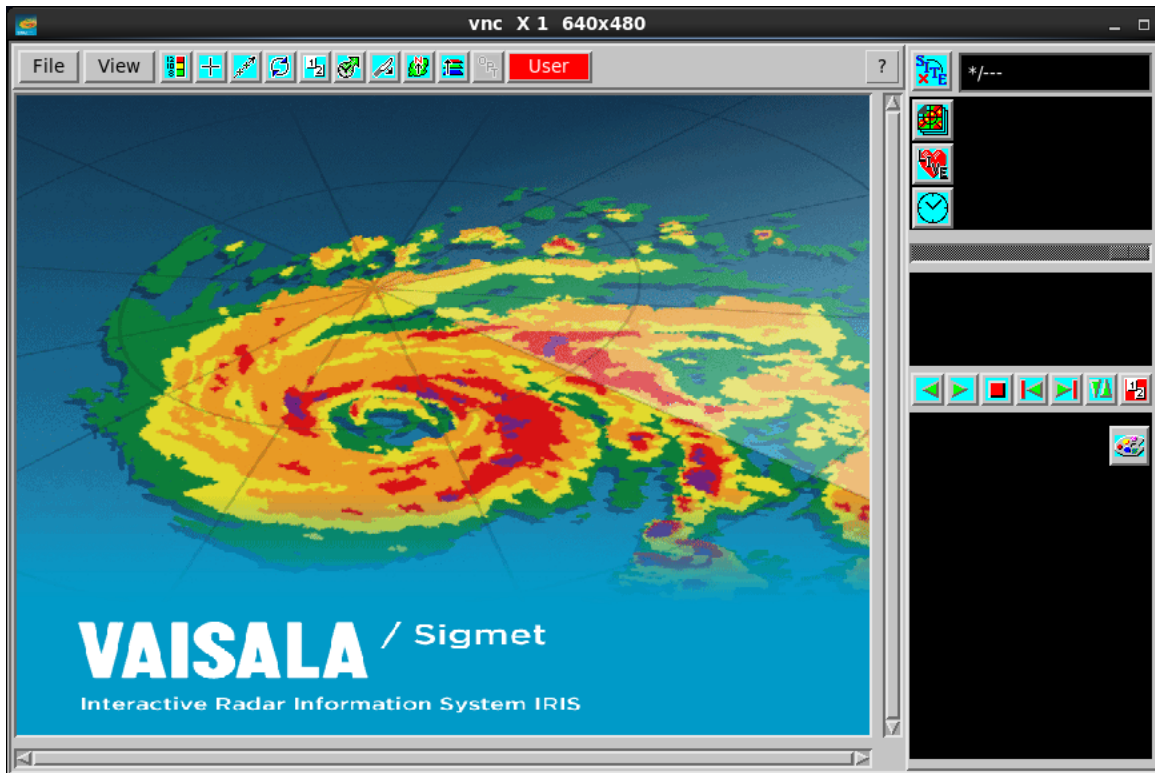
Figure 76   Quick Look Window

▶ 1. To select a site from all the sites configured at **General Setup**, select **Site**.
Select the site that corresponds to your local radar.

2. To display a list of available products, select **Product**.

   a. Typically at the radar site, you only have **PPI** products. To display **PPI** products,
   select a **PPI** from the list.
   If the list shows N/A, no products are available on the disk for the radar site you have
   selected.

   b. To display other products, select **PPI** to see all the available product types and then
   select the product.

   c. To display ingest data, in the upper right pane, select **Live** and a task name.

   RAW products cannot be displayed.

3. To display times for which there are images for the product name you have selected,
select **Time** (clock symbol).

4. Use the slide bar to browse older and newer images.

6.5 QLW Live Action Tool: Product Generation and Display

The **Live Action** tool allows you to configure, generate, and display many IRIS products. If one of these products is on the display, the Live IRIS tool becomes active.

Select the tool to show a light version of the relevant product configuration menu. Supported live action products include: **BASE**, **CAPPI**, **PPI**, **SHEAR**, **SRI**, **TOPS**, **VIL**, and **XSECT**.

Example: Live Action Tool Example for CAPPI
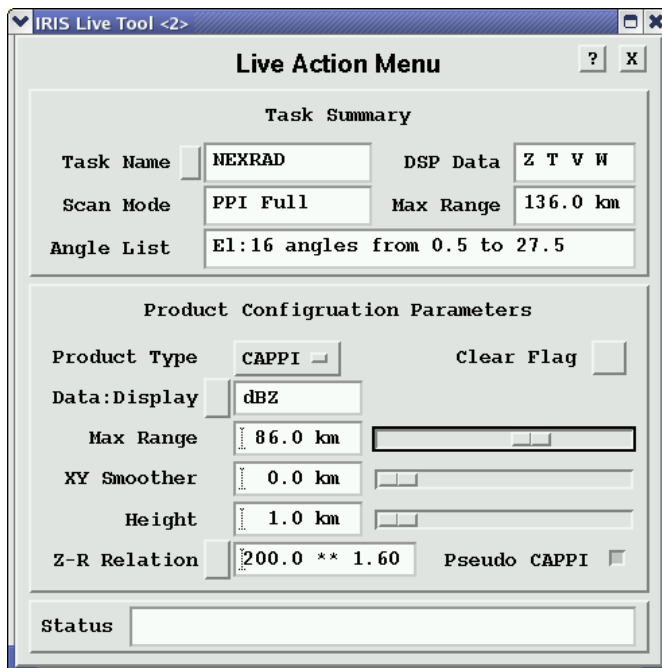
Figure 77   Example **Live Action** Tool for **CAPPI**

You can modify the product settings on the fly such as the cross-section line position or **CAPPI** height, and see continuous live updates on the display, typically at more than 10 frames per second.

For example, the **VIL** layer top can be scanned to observe the **VIL** at various heights above the freezing level.

You can scan the **CAPPI** hit scanned up and down to observe the vertical structure of the precipitation echoes. The **Live** menu also shows the type of display (Z, T, V, W).

Since the **Live Action** tool regenerates the product, it is necessary to send **RAW** data to the display so that the ingest files are available. Also, the display must be licensed for the **Live Action** tool.

More Information

‣ Configuring Products (page 17)

6.6 Changing Window Size and Zoom Level

▶ 1. On the top of the Quick Look Window, select **View**.

Zoom
1 X
2 X
3 X
4 X

Size	
Medium	480x480
Medium	640x480
Large	720x720
Large	880x720
XLarge	940x940
XLarge	1100x940

2. To change the zoom, select a zoom level and use the scroll bars to position your display in the area of interest.
   When you zoom, IRIS tries to preserve the center of the display.

3. To change the size, select a resolution.
    When you select a new window size, the image is rescaled to exactly fit in the new size. The sizes are the number of pixels in the image area for the window (not including the legend and boarders).
    *   Square image areas are recommended.
    *   Rectangular image sizes are useful if the legend information is included in the image itself (selectable as an output option in the QLW or in the **Product Output Menu** in the case of automatic outputs to a window).

> Try square images and use the Legend icon to toggle the legend on/off.

6.7 Color Scale Tool

The **Color Scale Tool** in the color legend lets you select different color scales that have been configured in your system in the **Color Setup** utility or create your own custom scale.

Figure 78   **Color Scale Tool**

1. To open the **Color Scale Tool**, select the **Color Scale** icon.

2. Choose the **Data Type** for which you want to define a color scale.
   Display Units shows the units of measurement for the data type that is currently selected.

   > When you first enter the **Color Scale Tool**, the data type does match what is displayed in the window.

3. Select a custom **Color Scale**.
   Custom scales are defined by the system operator in the **Color Setup** utility. There are up to 8 scales for each data type. These scales include the choice of color set, number of colors, and data levels for each color.
   The existing color scales remain intact in the software upgrades.
   When installing fresh software, a default color scale and the choice **Uniform** are available for each data type. In addition, the color scales **Winter** and **Everything** are pre-configured for the reflectivity data types, for the purpose of matching varied seasonal and application needs.
   The **Default** color scale of the `HClass` data corresponds to the class set of the legacy `HydroClass`. The `HClass` data type includes additional color scales `PRECIP`, `CELLS`, `METEO`, and `METEO-CELLS`, which correspond to the multiple Classifier methods (and their meaningful mergings) available in IRIS/RDA.
   The choice of **Uniform** has a special meaning: you can interactively select the color set, number of colors, and start and step levels. This is convenient for simple thresholding out low and/or high data values, as well as for the simplest mergings of echo classes of `HydroClass`.

4. If you are in uniform mode:

   a. Select a **Color Set**.
      Up to 4 named color sets are available. The color sets are defined by the system operator in the **Color Setup** utility.
      Typically you would define a **Default** set for positive valued data such as reflectivity; rainfall rate, **TOPS**, **VIL**, and so on. You would also define a **Velocity** set to be used for data that can be positive or negative such as velocity, ZDR or **SHEAR**.

   b. Type the number of colors to place in the scale (range 2 ... 16) in **Levels**.

   c. Type the start and step values in **First Level/Step**.
      The start value is placed in the upper edge of the first color. Control over saturation of the first and last color is taken from the **Default** color scale for the data type as configured in the **Color Setup** utility.

5. To highlight a data interval to see if that data is on your display, select **Highlight > Enable** and then **Edit** .
   The **Highlight Tool** starts. See 6.7.1 Highlight Tool (page 220).

6. Select **File > Save** to save your results.

> **i** The results are saved for only the window where you are working. You cannot affect other windows. You may want to go to the other windows and make similar changes and save them.

More Information

‣ VVP Output Options: Time-Height (page 258)

6.7.1 Highlight Tool

The **Highlight Tool**, launched from the **Color Scale Tool**, lets you highlight one particular data interval to see if that data is on your display.

Use this tool to select color intervals (from the 16 intervals in the legend). You can set the color for that interval using the color scroll bars. Those data values on the display are shown in the new color.

6.8 Display Options Tool - IRIS Product Output

The **Display Options** tool allows you to select display components such as geographic overlays, range rings, drop shadows, and even the overlay of multiple products, warning centroids, and echo tracks, along with adjusting the product overlay opacity. Select the icon to display the tool.
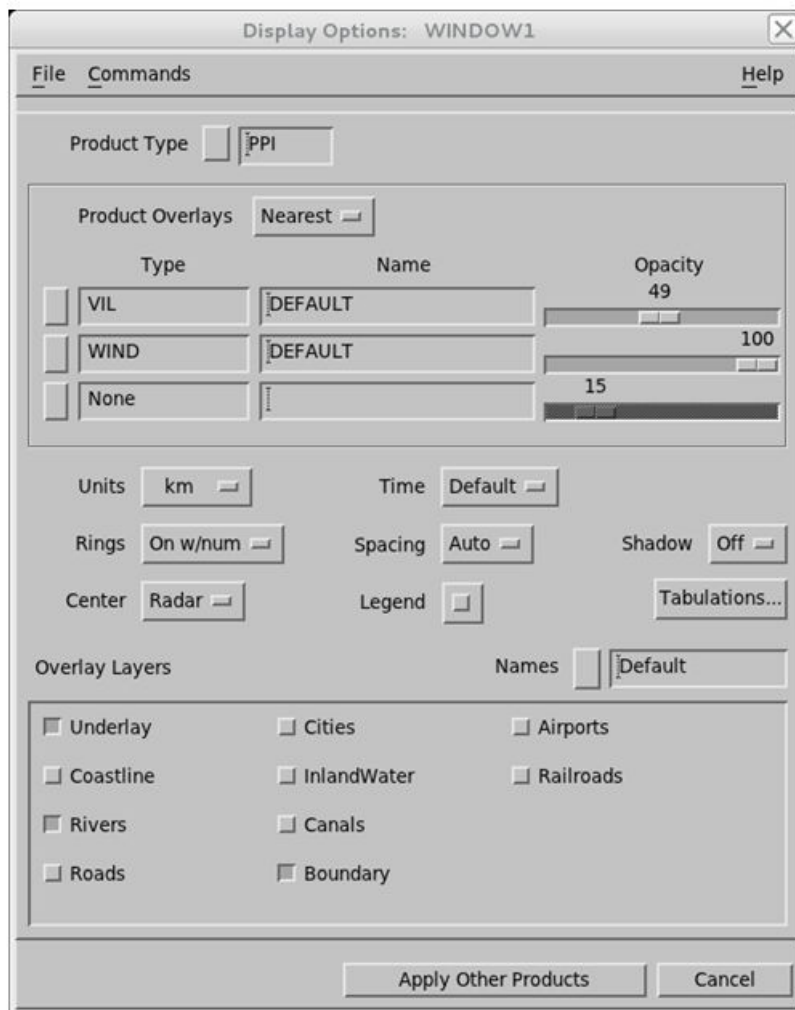
Figure 79   **Display Options** Menu

Saving Quick Look Window Display Options

Because there are many display options in this menu that can affect the appearance of a window, IRIS provides an easy way for you to save and recall configurations that you have made. Select **File** to show the following choices:

- **Save**
  Saves the menu current settings FOR THIS WINDOW ONLY.
  When you save a configuration you are not affecting any other window, any other user, or the operation of IRIS. This means that, with the permission of your system manager, you can make your own personal settings for a window without effecting anyone else.
- **Load**
  Restores the current saved settings. This is convenient if you want to make a temporary change.
- **Print**
  Makes a hardcopy of the menu to the default system printer.

Selecting the Product Type

When you first enter the **Display Options** tool the product type is inserted automatically in this field.

To choose a different product type, select the button.

To apply changes to this product type or all product types, use the apply options on the bottom of the window.

> If you display a **CAPPI** and change the product type field shown above to, for example, **PPI**, the changes are not applied are in the displayed product (the **CAPPI**) until you select **Apply All**.

Adding Product Overlays
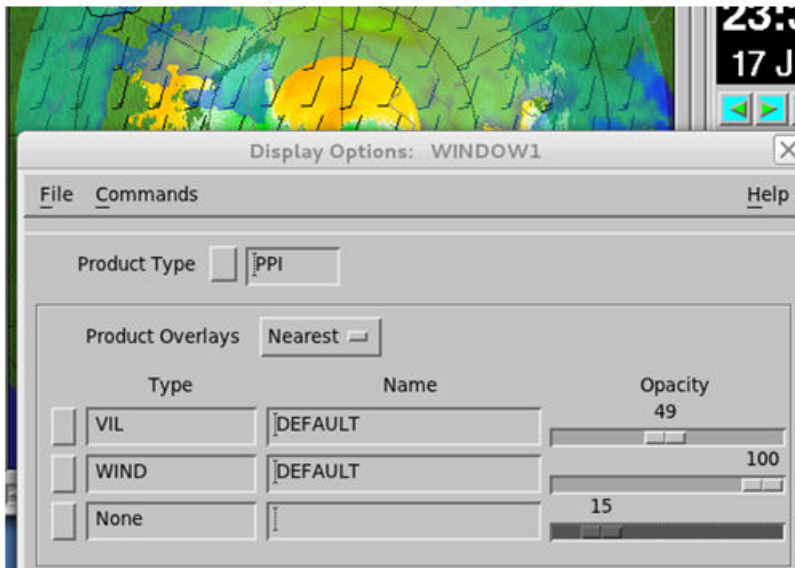
The **Display Option Tool > Product Overlay** section lets you combine different products. Almost any kind of product can be overlaid on any other product.

In the example above, an interactive **TRACK** product has been overlaid on a **PPI**. Typical applications for this are:

- **WARN** product overlays
- **WIND** or **NDOP** (multiple Doppler) wind barb overlays
- **TRACK**
- **SLINE** (shear line)

Multiple products can be overlaid in this manner. For example, **WIND** and **SLINE** displayed together show the wind vectors associated with a shear line.

Product overlays have time alignment options:

- **Nearest**
  The overlay product is the nearest in time to the base product.
- **Live**
  The overlay product is the most recent time.

It is also possible to combine 2 different types of image products. The example above shows **VIL** product with **TOPS** displayed underneath.

> When combining 2 image products, it is important to select contrasting color scales to be able to separate the data. In this example a gray scale was used for the **TOPS**. The drop shadows on the **VIL** are helpful.

> If your IRIS system supports satellite IMAGE product ingest, then this is how you combine the satellite display with radar data.

Range Units

This allows you to select range units for output displays. Choices are: scientific (kilometers) or navigational (nautical miles for horizontal distances and kilofeet for altitudes).

Time

This allows you to display UTC or the local time on the legend. The time has a Z suffix if it is UTC.

Select **Default** to let the **Setup** question to control this value.

Range Rings

You can add range rings, adjust the spacing of the range rings and center the rings on predefined points (for example, the radar or perhaps an airport. Select **Rings** to show the following:

- **None**
  No range rings are displayed.
- **On**
  Range rings are displayed but not numbered.
- **On W/Num**
  Range rings are displayed with range numbers.
- **Lat/Lon**
  A latitude/longitude grid is displayed. In addition, the display legend for each product shows the maximum product range.

Ring spacing is selected by the button to the right (**Auto** in the example).

Depending on the unit selection, the spacings are displayed as km or nautical miles.

The **Center** button allows you to select the center of the range rings from a list of pre-defined home points. The home points are defined in the **Cursor** or **Track** tool.

> You must have operator privilege to define home points in the **Cursor** or **Track** tools. Operator privilege for a **Quick Look Window** is set in the **File** tool by entering a password that is set in **Setup > Display**.

Legend

Enable this option to put a legend in the actual image. Most users leave this off since the **Quick Look Window** has a built-in legend.

If you elect to put a legend in the display, we recommend the following to optimize the window size:

- Use **View** to select a standard rectangular size (for example, 880 × 720).
- Use the Legend icon to toggle the **Quick Look Window** legend off.

Shadow

Enable **Shadow** to choose the kind of drop shadows displayed with products.

A drop shadow shades weather features, making them stand out against the background as if they were clouds shaded by the sun. You can turn drop shadows off, or choose from three shadow depths.

Tabulations

IRIS overlays can contain abbreviations or numbers for important places, such as airports, city centers, and similar.

Select **Tabulations** to show a menu with the full names for the abbreviations on the screen.

Overlay Layers
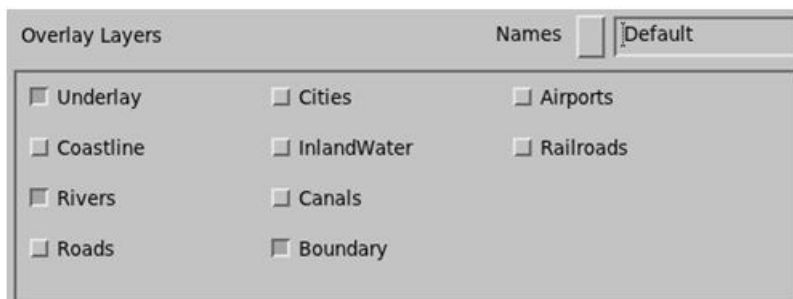
Geographic overlays show coastlines, roads, political boundaries and similar. See **Overlay** utility in *IRIS and RDA Utilities Guide*.

Select **Overlay Layers** to expand the **Display Options** menu to include the menu shown above. IRIS overlays are usually constructed in multiple layers so that users can turn on/off the layers.

After selecting some layers, select **Apply** to see your changes take effect.

"Virtual Overlays" are names given to different combinations of layers. To select a pre-defined virtual overlay, select the button next to the overlay name and select a new name from the list.

Operators can create "virtual overlay names." If you have operator privilege, enter your password through the **Quick Look Window** main **File** tool. Once you have privilege in the window:

1. Select the button next to the overlay name and select **New**.
2. Type the new name in the prompt.
3. Select the layers that you want.
4. To save your result, select **Display Options Menu > File > Save**.

When you save, the virtual overlay name is shared among all IRIS windows, printers, and so on.

Apply Other Products, Cancel

After you make changes to the **Display Options**, the **Apply Other Products** applies the changes.

Any changes that are applied affect only the window that you are in. You do not have to worry about affecting other windows (and vice versa).

Apply Other Products applies the changes to all product types, again for your window only. **Cancel** exits the menu without making any changes.

> **i** Press any button to apply your changes to the current product type for your window only.

6.9 Display Options Tool - Quick Look Window

The **Display Options** tool allows you to select components of a display such as geographic overlays, range rings, or drop shadows. Select the icon to display the tool.

Display Options: WINDOW1

File   Commands   Help

Product Type   [ ] [PPI]

Units   [km ⌐]   Time [Default ⌐]

Rings   [On w/num ⌐]   Spacing [Auto ⌐]   Shadow [Off ⌐]

Center [Radar ⌐]   Legend [ ]   [Tabulations...]

Overlay Layers   Names [ ] [Default]

☐ Underlay   ☐ Cities   ☐ Airports
☐ Coastline   ☐ InlandWater   ☐ Railroads
☐ Rivers   ☐ Canals
☐ Roads   ☐ Boundary

[Apply Other Products]   [Cancel]

6.10 Product Overlay Tool

The **Product Overlay** tool allows you to overlay multiple products, warning centroids, and echo tracks, or adjust the product overlay opacity.
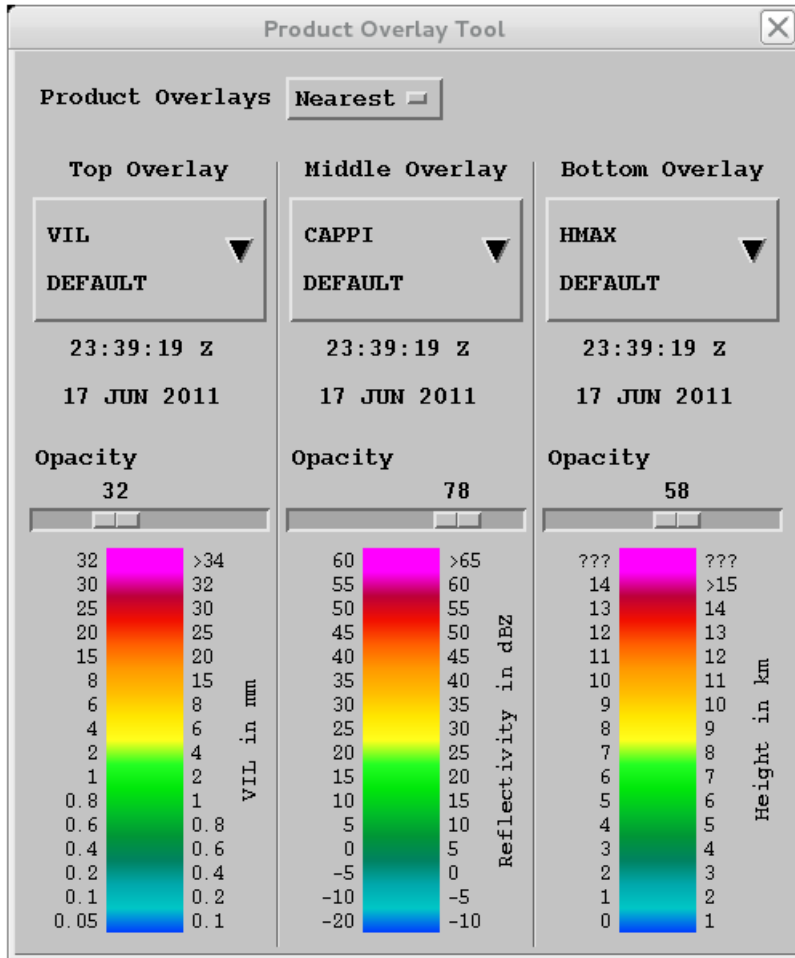
Figure 80   **Product Overlay** Tool

▶   1.   Select the icon to show the **Product Overlay** tool.

2.   Select the product type and product name.

6.11 Animation Tool

Animation is one of the best ways for forecasters to observe the motion, growth, and decay of radar echoes. In real time applications, many IRIS users leave a loop running constantly.

The **Animation** tool lets you control the animation sequence.

> You can also use the **Slide Show** tool to automate the display of different products from different times.
> See 6.12 Slide Show Tool (page 234).

Loop Start and Stop

Use these buttons to play (reverse and forward play) and stop animations. They are also available in the **Legend**.

> To use the SPACE BAR to stop a loop or start it in the forward play direction, place the cursor in the image area of the loop.

> When a loop starts, the Site filter automatically matches the loop so that only images from the site (for which the loop was started) are displayed.

Loop Loading and Status

When a loop loads, it shows how many frames have been loaded and how many total frames are available (29/30 in the example). The loop starts playing immediately, while it loads. You do not have to wait for the loop to be fully loaded before it is playing at loop speed.

After the loading is complete, a number, 30 in the example, indicates the number of frames in the loop. The time span shown is the actual span of the frames in the loop, not the requested **Max Time**.

Factors that determine the number of frames in a loop include:

- **Device Movie Length** and default window size configured in **Setup > Output**.
  If your window size is the same as the **Setup > Output** default size, this is the maximum number of frames shown.
  If you set your window size to be smaller than the default size, then you are shown more frames, since the same amount of memory is allocated for movies.
  If you set your window size to be larger than the default, you get fewer frames.
- **Max Time** determines the maximum possible span of the movie. This limits the number of frames.
- The number of frames available on disk. To see this, check the **Time** icon.

Flush the Loop Buffer

The loops are "sticky". The frames are kept in a loop buffer so the loop does not have to reload if it stops and restarts. This is convenient when you are working with long loops.

If you make a change to the window such as the color scale or the **Display Options**, the buffer automatically "flushes" and the loop reloads your changes.

If your loop must be flushed and reloaded or your loop buffer becomes corrupted, use the **Flush** button to refresh the buffer. The loop reloads the next time it starts.

Previous/Next

Single step forward and back through the loop. This is similar to using the **Time Slider** or using the left/right arrow keys on the keyboard (with your cursor in the image area).

Toggle

Swaps between the current image and the previous image on the display, regardless of the image type or time. Sometimes this is useful for comparing 2 different products (for example, **TOPS** and **VIL**).

Display both products, then use the **Toggle** icon to alternately display each one.

> It is convenient that when the cursor is used, the position is displayed correctly in the 2 toggle frames, even if their scales are radically different.

Skip Time

If you want to loop a long time series of data, for example, say 24 hours, you must usually necessary skip some frames so that you can display the entire 24-hour span without exceeding your `Device Movie Length` for the maximum number of frames.

Specify the time skip in `HH:MM`. For example, specify 01:00 (one hour), for a 24 hour movie to have 24 frames.

The skipping algorithm attempts to show the frames that are exactly on or as close as possible to intervals of the skip time. In the case of a one hour skip time, the frames would be on or as close as possible to just after the hour.

Max Time

This is the maximum allowed time span.

You may not be able to achieve this since there may be insufficient frames on the disk or you may be limited by the `Device Movie Length` frame/memory limit configured in **Setup > Output**. The actual achieved time span is displayed in the **Span** field.

Speed

This is the requested speed in frames per second. Faster is to the right, slower is to the left.

Achieving the speed depends on the speed of your CPU, the amount of memory that you have, and the size of image that you are looping.

Pause

This is the pause in seconds between successive loops.

In addition to marking the end of a loop, **Pause** gives the computer a chance to rest and do other things. For example, setting a very short pause of 0.0 sec runs the loop continuously, which can cause other tasks such as loading the loop to take longer. If your computer seems too busy, lengthen the pause. The 1.5 second default value is a good compromise.

> ⚠️ **CAUTION!** High-speed loops of large images place a high demand on CPU and memory.
> By running several large format loops at high speed with no pause, you can use near 100 % of your CPU and memory depending on your system hardware configuration. Consider the IRIS tasks on your system such as product generation and data acquisition from the radar. The default values for **Speed** and **Pause** are reasonable starting points.

Setting the Loop End Time
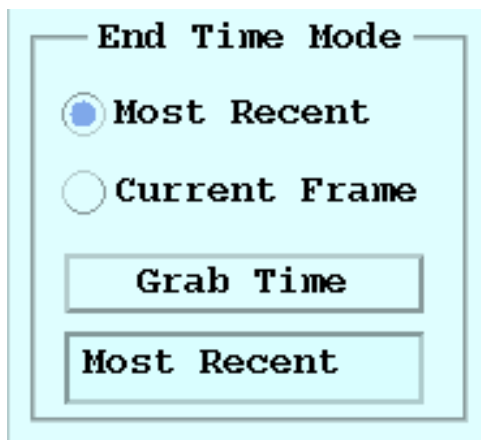
Usually the last frame in the loop is set to be the **Most Recent** picture. In this case, new data are added automatically to the loop. This allows you to leave a loop running on a display so that you can always see the current situation. Note that when the loop end time is set to most recent (the default), the **Mode** indicator at the top shows:

The green background behind **User** indicates that the display automatically updates with new images. If the loop does not reflect the most recent data, the **Mode** shows a red background color.

The **Current Frame** and **Grab Time** options allow you to fix the end time for occasions when you want to create loops that do not end with the most recent image, for example, for analysis of archive data from last year.

- **Current Frame** mode sets the loop end time to be the time of the image that is currently in the window when the loop is started. Browse through the images until you decide what frame should be the last frame and then start the loop. The loop ends on the frame that you selected.
  As you start and stop loops in this mode, the end time is not fixed, it is readjusted every time that the loop is restarted.
- **Grab Time** allows you to select a specific time for the end of the loop. Browse through the images until you find the image that you want to be the last frame and select **Grab Time** to latch this time. Even if you change product types, any loops started use this as the end time.

> ℹ **Current Frame** and **Grab Time** the loops do not update automatically.
> Avoid leaving the displays in these modes since it could confuse other users.

6.12 Slide Show Tool

You can use the **Slide Show Tool** to observe the motion, growth, and decay of radar echoes.

The **Slide Show Tool** can display different products from different times (while the animation tool is limited to a linear time sequence of one product). For example, in the following figure, different 4 slides of different products are shown in the Quick Look Window (QLW) for 10 seconds.

You can configure multiple windows so that images in different windows change at the same time.

> ℹ For multiple window presentations it is convenient to use **Setup > Output** to configure the start-up window location, window size, and legend on/off.

> 💡 You can use **Slide Show Tool** for a weather briefing display for which there is no operator/observer control of the display.
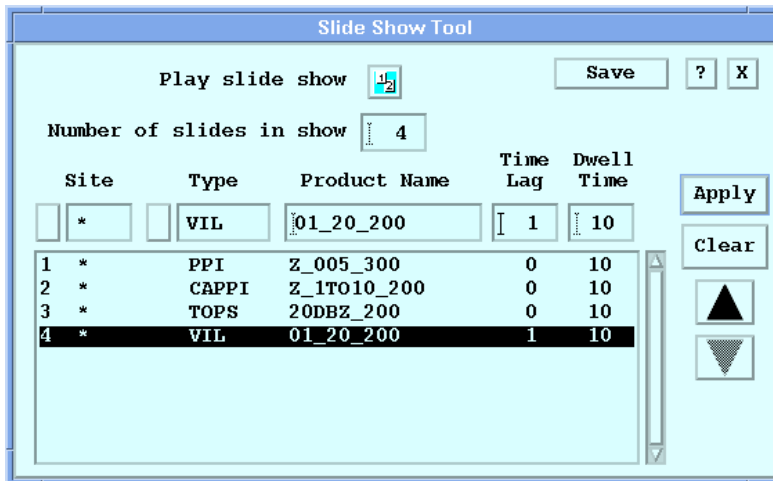
Figure 82   **Slide Show Tool**

▶   1. Select the slide show icon to display the tool.

2. Enter the number of slides in **Number of slides in show**.
This field determines the number of slides in a show (maximum 16).
IRIS inserts the specified number of blank rows in the slide list.
You must populate each slide with information about the product it will display.

3. For each slide select the slide number (or row) and do the following:

   a. Specify the site, type, and product name.
   Select from products on the disk.
   If there is no product of the type that you want, for example, if it has not been made yet, select the blank product name "––––––" and type the name you want. When the product has been transferred to your disk, it is displayed in the slide show.

   b. Specify the **Time Lag**.
   Time Lag specifies the version of the product according to the following rules:
   - **0** The current (most recent) version of the product.
   - **1** The previous version of the product, that is, the second most recent.
   - **2** Two products prior to the current version.

   This enables the comparison of product versions. In a multi-window environment, the newest version of a **CAPPI** product can be displayed alongside a **CAPPI** product from the previous hour.

   > To make a single step time sequence, type the same product type, the name for each slide, and specify 0, 1, 2, … for the lags.

   c. Specify the **Dwell Time**, that is, how long the slide is displayed on the screen, before sequencing to the next slide.
   In the example menu, each slide is displayed for 10 seconds.
   Dwell Time defines how long a slide is on the screen as well as a precise clock schedule of time slots for displaying the slides, that is:
   - Slide 1 Displayed at Midnight 00:00:00
   - Slide 2 Displayed at 00:00:00 plus the Slide One Dwell Time.
   - Slide 3 Displayed at 00:00:00 plus the Slide 1 and Slide 2 Dwell Times

   The schedule increments throughout the day.
   Note that a slide show can start at any time, but the time slots are defined by the above algorithm. This allows you to configure multi-window slide shows where the images in different windows all change at the same time and in the proper sequence.

   d. Select **Apply**.

4. To clear incorrect slide information, select the slide number (or row) and then **Clear**.

5. Select **Save**.

6. To start the show, select **Play slide show** or select the slide show icon in the legend.

6.13 Cursor Tool

You can use the **Cursor** tool to show the value of data at a point and the geographic position of the point. In addition, the position of the cursor relative to another "home" point (range and bearing) are displayed.

You can select the "home" point from a pre-defined list that always includes the radar, or you can "plant" it interactively.
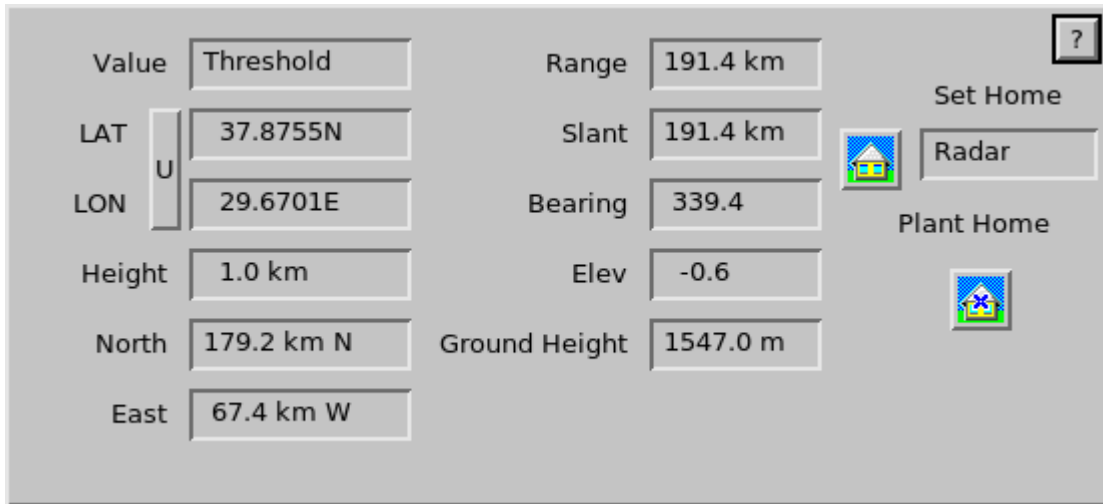
Figure 83   **Cursor** Tool

Value

 Shows the value at the current position. The units are the same as those displayed in the image.

Latitude and **Longitude**

 Show the latitude and longitude of the position to the nearest tenth of a minute. Select **U** to toggle the units between decimal degrees and decimal minutes.

Height

 Shows the height above the curved surface of the Earth at the cursor position. [1]
 Not all displays have a height, **VIL** for example. In these cases, a value of `0.0` is displayed for height.

North

 Shows the North–South distance from the "home" point to the cursor position. [1]

East

 Shows the East–West distance from the "home" point to the cursor position. [1]

---

1) *The units are either km or nautical miles depending on the setting of* `Units` *in the **Display Options** tool. See in 6.8 Display Options Tool - IRIS Product Output (page 220).*

Range and **Slant** Range

Range shows the range along the curved Earth's surface relative to the "home" point (usually set to be the radar). [1]

Slant shows the range along the radar beam. [1]

Not all displays have an associated **Slant** range because there may be no elevation angle associated with the product. **VIL** is a good example. In these cases, the **Slant** is displayed as `0.0`.

Bearing

The bearing in degrees from the Home point to the cursor as measured clockwise from true north.

Elevation

For products that have an associated elevation, this displays the elevation angle from the home point to the cursor.

Any product that has a "height" also has an elevation (for example, **PPI** and **CAPPI**). If the "home" point is relative to the radar, the elevation is the elevation angle of the radar.

If the "home" point is not the radar, the elevation is the elevation angle that an observer at that point would observe to see the cursor at its height.

Ground Height

Displays the height of the ground above sea level at the location of the cursor.

To use **Ground Height**, you must add Digital Elevation Model (DEM) data in GeoTiff format in the */usr/sigmet/dt* directory with the filename *dem.tif*. The DEM data should cover the geographical domain scanned by radars within the network. Good quality DEM data in GeoTiff format is available at, for example: http://srtm.csi.cgiar.org.

> You can use cursor with other tools such as the **Loop** or **Track** tools. However in some cases, the mouse is reserved for other functions.
> You can still get readouts from the cursor tool by displaying it first and then selecting **Track**. The cursor is not displayed on the screen, but the readouts correctly reflect the values where you are pointing.

1. Select the Cursor icon to display the tool.
   You can move the cursor in the display by clicking your mouse, or dragging the cursor. The values in the menu update as you move the cursor.

6.13.1 Selecting and Planting Home Points

You can define home points () to measure the position of a weather radar feature relative to a selected location.

For example, `Thunderstorm located 120 km at 260° from the airport.`
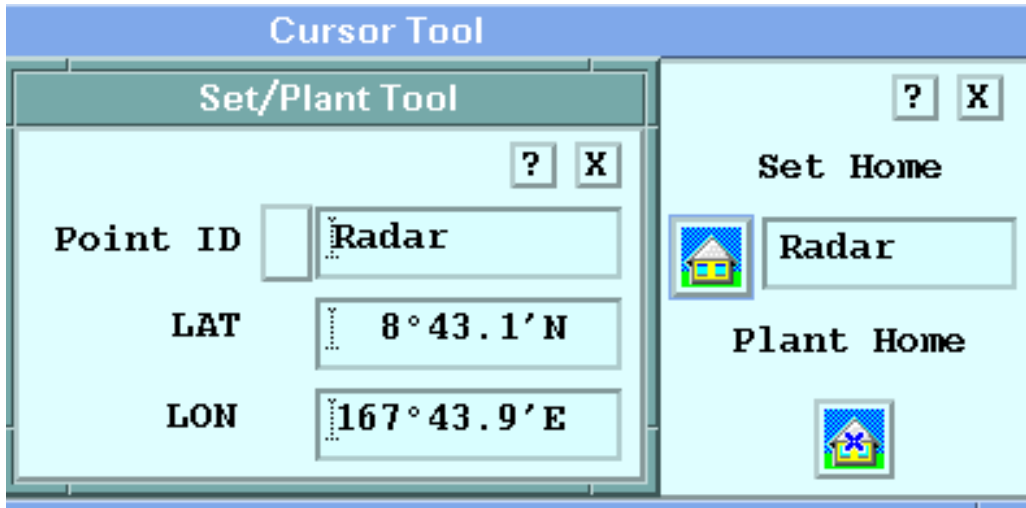
Figure 84   Cursor **Set/Plant Tool**

1. In the **Cursor** tool, define home points by doing one of the following:
   - Use your mouse to position the cursor and select **Plant Home**.
     The Home point, indicated by an **X** moves to the cursor.
   - Select the **Set Home** icon to show the **Set/Plant Tool**.
     Select a **Point ID** and then the **Home** point from the list.

6.13.2 Creating Named Home Points (Operators only)

If you have operator privilege in the window (**File > Set Privilege to**), you can use the **Set/ Plant** tool to make new named points for key locations.

1. Select . In the **Set/Plant**, do one of the following:
   - Enter the LAT/LON of the home point and the name of the home point. You can type the LAT/LON using either degrees and minutes or decimal degrees.
   - Select **Point ID** and select **Save**.

2. Position your cursor where you want to place the home point and then click the **Plant Home** icon.

   a. Select **Set Home** to fill in LAT/LON automatically.

   b. Replace the text `Plant` with the name you want.

   c. Select **Point ID** and select **Save**.

3. To delete a home point name:

   a. Using **Point ID** to select the point.

   b. In the **Point ID** menu, select **Delete**.

   The Radar is used as the home point until you select another one.

6.14 Track/Annotate Tool
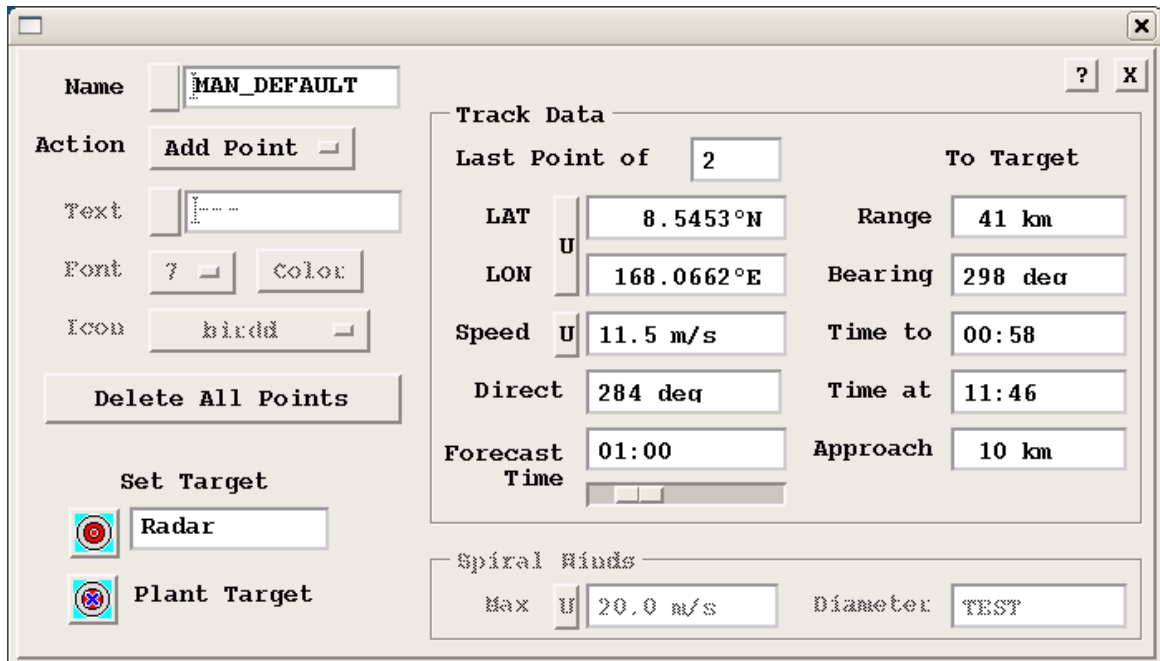
Figure 85   QLW Track/Annotate Tool

The Track/Annotate tool makes it easy for forecasters to manually:

- Make and edit interactive storm tracks.
- Determine time of arrival, time to arrival and distance of closest approach to a selectable "target" point (similar to a home point in the cursor mode).
- Add annotation text in various font sizes and colors.
- Add bitmap icons.
- Determine the center of a typhoon spiral.

The left side of the menu contains the control of the next the cursor action and selection of the forecast target point.

The right side of the menu contains the track and forecast information, such as the position of the most recent track point and the distance, bearing and time of arrival from the most recent track point to the selected target point.

Example of Track Points, Icons and Annotation

The elements of a track are shown in the example above.

- Track Points are indicated by the "+" with the time displayed beneath each one. Note that the latest track point at 1926 is highlighted since this corresponds to the frame that is currently displayed.
- **Forecast Line** or arrow shows the expected position of the echo, in this case, in 1 hour. This line is also displayed in the highlight color to distinguish it from the track line.
- **Icon** in this case is an airport symbol.
- Annotation examples are the text "Icon", "Target" and "Forecast Line".
- **Target** for the forecast is indicated by the "X" symbol. This is used for forecasting the time of arrival at a position.

 Target Information

The "Target" is a geographical position that you use for forecasting purposes such as an airport or urban area. The default target position is at the radar site, but you can select the target from a list or plant a target with your cursor, similar to home points in the cursor menu.

The list of Target points is the same as the list of Home points. These can only be saved or deleted by the operator. See 6.13 Cursor Tool (page 236).

Once you have set your target, the target information shows:

Range and **Bearing**
> To target from the most recent track point. In the example, the target bears 118° and 19 km.

Time to (target)
> The time to the closest approach of the track to the target.

Time at
> The clock time of when the echo makes its closest approach.

Approach
> The distance of closest approach as measured from the track line to the target.

6.14.1 Creating and Editing a Track

▶ 1. To start a new track:

   a. Select an existing track **Name** or type a new name.

   b. Select **Delete All Points** to start a fresh track.
      The menu shows `Last Point of 0` to indicate that all points have been removed and the **Action** shows **Add Point**.

2. In **Name**, type the name of your **TRACK** product.

The Track **Name** field shows the name of a manual **TRACK** product.
The names are always preceded by `MAN_` to indicate that these are manual tracks rather than automatic **TRACK** products created by IRIS.

- To select an existing track name, select the button next to the name field and select from the list of pre-defined names. There is always a track with the name `MAN_DEFAULT`.
- To create a new name, replace the text in the name field. IRIS automatically adds the `MAN_` prefix and converts the name to upper case.
- To delete a name, select the name you want to delete and select **Delete**. The name is deleted from the list and the `MAN_DEFAULT` is inserted as the active track.

3. In **Track Data**, define your track.

```
┌─ Track Data ──────────────────────────────────────┐
│  Last Point of    ┌─────┐           To Target      │
│                   │  2  │                          │
│                   └─────┘                          │
│     LAT        ┌──────────────┐   Range  ┌────────┐│
│            ┌─┐ │   8.5453°N   │          │ 41 km  ││
│            │U│ └──────────────┘          └────────┘│
│     LON    │ │ ┌──────────────┐  Bearing ┌────────┐│
│            └─┘ │  168.0662°E  │          │298 deg ││
│                └──────────────┘          └────────┘│
│   Speed ┌─┐┌──────────────┐   Time to ┌──────────┐ │
│         │U││  11.5 m/s    │           │  00:58   │ │
│         └─┘└──────────────┘           └──────────┘ │
│   Direct  ┌──────────────┐    Time at ┌──────────┐ │
│           │  284 deg     │            │  11:46   │ │
│           └──────────────┘            └──────────┘ │
│  Forecast ┌──────────────┐   Approach ┌──────────┐ │
│    Time   │  01:00       │            │  10 km   │ │
│           └──────────────┘            └──────────┘ │
│           ┌──────────────┐                         │
│           │   ▭▭         │                         │
│           └──────────────┘                         │
└────────────────────────────────────────────────────┘
```

The center column of the Track tool shows information about the track you are constructing.

Last Point of <N>

N shows the number of track, annotation and icon points that have been defined.

LAT/LON

Shows the latitude and longitude of the last track point, that is, the most recent data time.

Select **U** to toggle the units between decimal degrees and decimal minutes.

Speed and **Direction**

Shows the speed and direction of the weather feature as computed from the two most recent track points.

Select **U** to select common units.

Forecast Time

Allows you to set the time of the forecast.

4. To add points to a new track:

    a. Select a product for tracking.
        For example, if you want to track the most severe thunderstorms, choose **VIL** since it is a good indicator of storm severity and tends to be persistent.

    b. Browse through the images to select a feature, place and click the cursor on the feature.
        A track point appears along with text to indicate the time of the point.
        The color of the track point is in a special highlight color to indicate that it corresponds to the displayed image.

    c. If you do not like the position of the track point, click again to move the cursor to a new position.
        Drag the point to where you want it, and release.

    d. Continue browsing the images and adding track points.

> To avoid crowding the track points, skip a few images. You do not need a point for every image.

> In track mode, use the mouse to add track points and the left/right arrow keys for browsing.

5. To delete a track point:

    a. Select **Action > Delete**.

    b. Select the track point that you want to remove.

> Keep pressing the mouse button to highlight the point to be deleted. This is helpful when you want to delete a track point that is close to other track points.

6.14.2 Adding and Editing Annotation Text

1. Select **Action > Add Text**.

2. Type the text that you want or select an existing saved text string.

3. Select the font and the color of the text.

4. Position the cursor where you want the text to appear and click.
    The cursor shows you a few preview characters in the selected font and color so you can decide before you click if the text is OK.

5. To delete, edit, or move a text string, select **Action > Delete**.

6.14.3 Adding and Deleting Tracking and Annotation Icons

▶ 1. To add an icon such as an airport or thunderstorm symbol:

   a. Select **Action > Add Icon**.

   b. Use the **Icon** button to select the type of icon.
      Usually the name suggests what the icon is.

   c. Use **Color selection** to select the icon color.

   d. Position the cursor (now shaped like your icon) on the screen and click.

2. To delete an icon:

   a. Select **Action > Delete**

   b. Select the icon to be removed.

3. To create your own icons, use a program such as Paint (on a PC) or Kiconedit (on a Linux system).
   Icons are *.xbm* format bit-mapped icons which are stored in the */usr/ sigmet/ config/overlay* directory.
   To get you started, Vaisala provides several sample icons.

6.14.4 Tracking Typhoons With the Add Spiral Option

Use this mode to determine the location of a typhoon center, and to calculate the maximum rotational winds. Even if the center is out of range of the radar scan, you may be able to line-up the spiral line with the observed data to determine the center.

Figure 86   Typhoon Track Example

The elements of a typhoon track are shown in the example.

- Spiral Band Line is drawn on the display. You can drag the center point to change the location, or drag the end of the spiral to change the rotation. The goal is to line up the drawn spiral with one of the rain bands.
- Typhoon Track Points are indicated by the "+" with the time displayed beneath each one. Note that the latest track point at 1048 is highlighted since this corresponds to the frame that is currently displayed.
- Forecast Line or arrow shows the expected position of the eye, in this case, in 1 hour. This line is also displayed in the highlight color to distinguish it from the track line.

- Target for the forecast is indicated by the "X" symbol. This is used for forecasting the time of arrival at a position.

▶ 1. In the **Track/Annotate > Spiral Winds** section, calculate the maximum winds.

The IRIS QLW wind model looks at the data and finds the best fit of a rotation about the center point, plus a mean horizontal translation.

a. Enter the values.
   - The maximum wind and the wind speed at that diameter.
   - The average translation speed and direction.

b. Select **Compute**.
   The text box to the right of the **Compute** button shows the status of this calculation. Note that this calculation requires a lot of CPU.
   You must have the polar velocity data on your serving machine. If this data is unavailable, the **Compute** option is grayed out.

If you move the spiral location, the status displays `Repositioning Spiral`. This remains until you push the button. In this case, it displays `Fitting winds`, and finally `Success`. It can take up to a minute to do the fit.
There are a number of possible errors in the fit. For example:
- If the weather is clear and there are no velocities in the data, the status shows `Error: 0 velocity points`.
- If the fit cannot be solved, the status shows `Error: Bad fit`.

If you try to fit a cyclonic motion to a data set which is not rotating, you get unuseful numbers such as a small rotation speed at a large range. Also because the radar cannot detect rotation about the radar, if you place the cursor within a few km of the radar, the numbers are invalid.

2. To change the tightness of the spiral line, use the **Angle** widget, and increment the angle up and down from 5 to 20°.

This is the angle in degrees that the logarithmic spiral makes to a circle. A smaller number makes a tighter spiral.

a. Set the direction of cyclonic motion at your location to **CW** or **CCW** using the **Rotate** option.

b. Set the minimum and maximum displayed range of the spiral using the **Range** type in fields.

6.15 Forecast Tool

Use the **Forecast Tool** to shift products into the future to see the forecast position of weather features.

In the following example, the dark outer region represents the area not scanned, while the inner lighter circle represents the scanned area. The radar image is shifted 3 hours for an assumed motion of 24 knots at 43°.
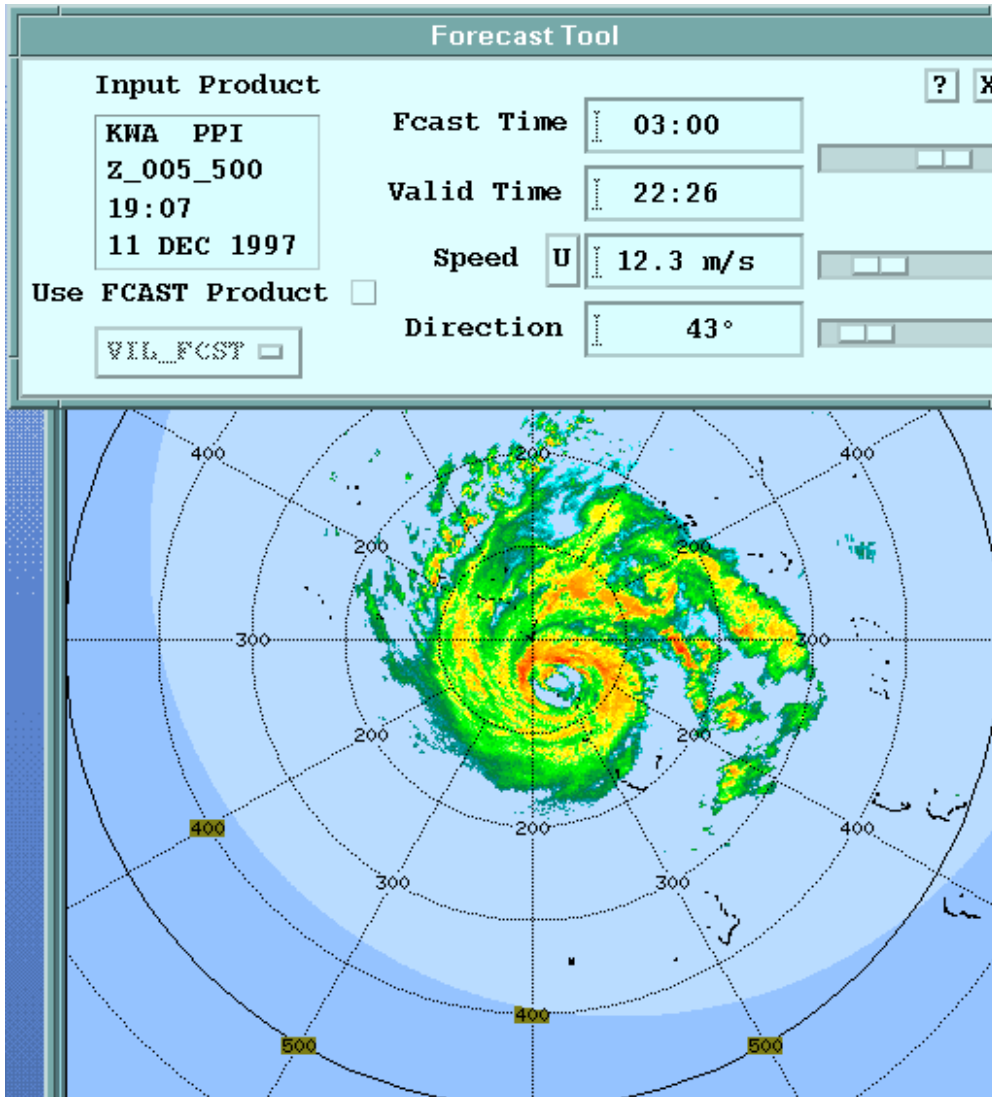
Figure 87   Forecast Example with a Typhoon

The image can be shifted in the following ways:

- User Input
  The forecaster inputs a speed and direction as was done in the example.
  The echo is shifted by a constant speed and direction. This is shown in the example above.
- **FCAST** product (IRIS or user Forecast product)
  Forecast products can be produced either by IRIS or by external programs that create forecast products and insert them into IRIS.
  The echoes can be subjected to both a non-uniform motion field and a non-uniform intensity change field.
  The IRIS **FCAST** product computes a uniform wind field based on the average echo motion as calculated from a 2D cross-correlation of the current and previous image. No intensity change is calculated. Check with your system manager to see what **FCAST** products are available on your system.

User Input Forecast

If you have obtained a nominal speed and direction from the manual track product, you can do a manual forecast projection with the following settings:

- **Use FCAST Product** button clicked out.
- **Speed** and **Direction** are set manually by type-in or slide pot.
- **FCAST Time** is set to how far ahead you want to look.
- **Valid Time** displays the time at which the forecast is valid, that is, the data time of the display plus the **Forecast Time**

FCAST Product Forecast

- Select **Use FCAST Product**.
- Select the forecast product, if any are available, from the option widget.

In this case the speed and direction are taken from the forecast product and the user input fields are desensitized.

The **FCAST** algorithm in the **Quick Look Window** takes as input a product and a forecast product (array of motion vectors and intensity changes).

The steps involved in the forecast are as follows:

1. An output array is created with the same projection and resolution as the input product (for example, an input **CAPPI** that we want to forecast).
2. The forecast product motion vectors and intensity change array points are re-mapped onto the output array. This re-mapping means that the input product and the forecast product do not have to have the same projection, resolution, or range.
3. Usually there are many more output array pixels than there are grid points in the forecast product. Therefore it is necessary to interpolate the spares motion vector and intensity change points to each pixel point (pixel) in the output array. The result is that each pixel of the output array has a motion vector and an intensity change associated with it.
4. For each pixel of the output array, the motion vector is used to "look back" in time per the selected forecast time. This "backwards displacement" vector points to a pixel in the adjusted input product. The value of this pixel is then used for the output product pixel.
5. The intensity changes in the output array are applied pixel-by-pixel to adjust the data for growth and decay.

More Information

‣ FCAST: Forecast (page 40)

6.16 Cross-Section (**XSECT**) Tool

The interactive cross-section is a good way to see storm structure and echo heights.
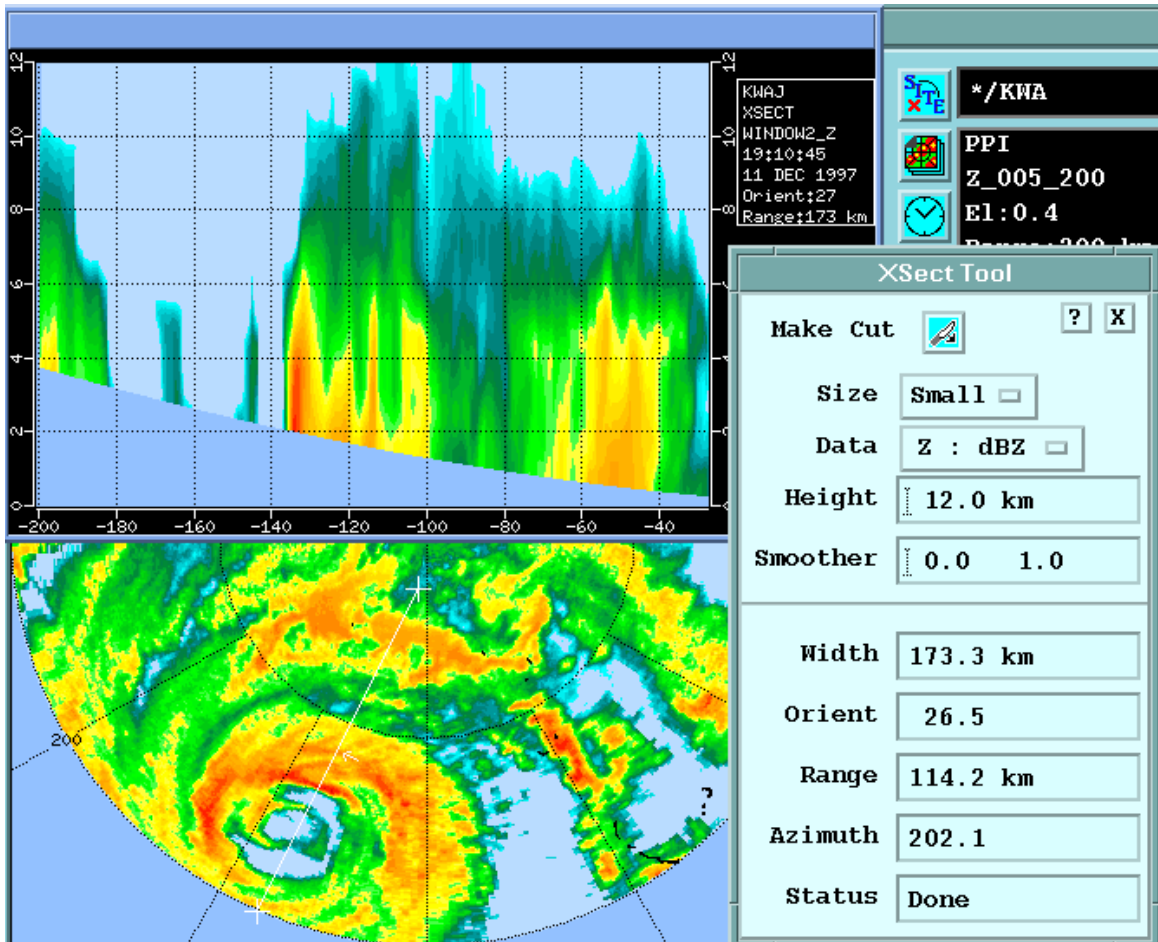
Figure 88   Cross-section with Typhoon Example

The window at the top left is a small format **XSECT** window. This appears automatically when you make a cross-section. Note that because the radar cannot see all the way to the surface of the Earth, the bottom of the cross-section is a curved boundary.

The cross-section line is shown on the display. The small arrow at the center of the line shows how you see the cross-section. In this case the cross-section is viewed from the east through the eye of the typhoon.

The **XSECT** tool is shown on the right of the example allows you to configure nearly all aspects of the cross-section.

Cross-Section Display Color Scales

The cross-section uses the color scale configured in the **Color Scale** tool for the data type that is being displayed.

The cross-section and base product color scales match if they have the same data type. For example, if you ask for a cross-section of **dBZ** and the base product is **dBZ**, the color scale of the cross-section match the input product.

If you want to change the color scale in the cross-section, use the **Color Scale** tool and select the data type (`dBZ`, `V`, and similar). Then re-make the **XSECT**. The new colors are applied when the cross-section is re-made.

Cross-Section Range Scale

The height range scale is either km or thousands of feet and the horizontal range scale is either km or nautical miles according to the units selection in the **Display Options** tool.

The origin (range zero) corresponds to the point on the line or its extension that is closest to the radar. That is, the origin is the point where the line or an extension of the line would be at a right angle to a line drawn from the radar.
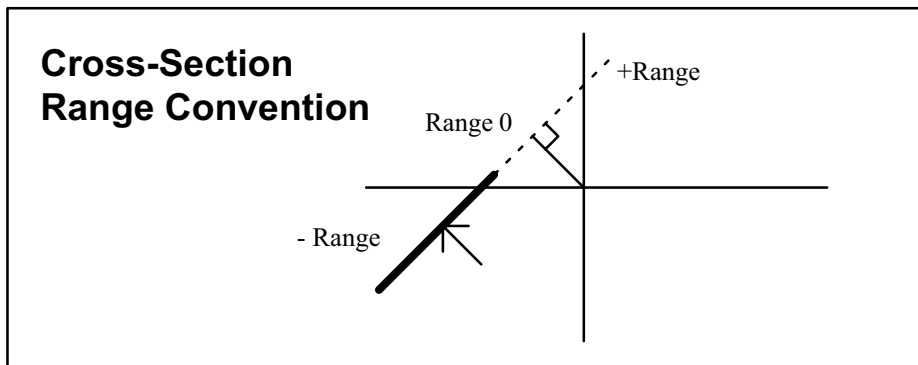
Figure 89   Cross-Section Range Scale

Cross-Section Display in Standard Window

To display ccross-sections in a standard window select an **XSECT** product.

The cross-sections inherit the name of your window and the data type when they are listed in the product selection tool.

6.16.1 Creating a Cross-Section

> Cross-sections are generated by interpolating original polar data.
> To make a cross-section, you must have ingest files from the signal processor or from RAW data on your system. If the window is exported from another system to your screen, that system must have the ingest files.

1. Select an appropriate "base" product.
   You cannot use products such as **RAIN1** and **RAINN** for a cross-section since they do not represent a specific time.

2. Select the **XSECT** tool (the icon with the cutting knife).

3. Use the mouse to position the endpoints of the cross-section line.

4. Configure the cross-section:

    a. In **Height**, specify the height of the cross-section in km to the nearest 0.1 km (100 m) (328 ft 1 in).

    Height ┊ 12.0 km

    b. In **Data**, select the data types from the list.
    The left variable represents the input data type, the right variable represents what is displayed.
    The following example shows a polarization diversity example.

    Data   Z : dBZ ▢

    | Z : dBZ |
    | Z : R |
    | V : V |
    | W : W |
    | T : dBZt |
    | T : R |
    | Zd : ZDR |
    | Kd : Kdp |
    | Kd : R |

    c. In **Size**, select a large or a small size window.

    Size   Small ▢

    A small size is usually better since it takes up less screen area and generates more quickly.

    d. In **Smoother**, select the horizontal and vertical smoother length scale in km.

    Smoother ┊ 0.0   1.0

5. To cross-section, select the **XSECT** icon.

6. Monitor the cross-section line configuration progress in the **Status** field at the bottom of the tool.
   - **Range** and **Azimuth** refer to the center point of the line.
   - **Status** shows the progress: `Setup`, `Running`, and `Done`.

   If the ingest data for the requested data time (that is, the time of the base product) are not on disk, the status shows `Error` and a message informs you the ingest files were unavailable. In this case, browse forward in time and try again, or check the **Ingest Summary Menu** to see what data are available.

6.17 Product Output Options Tool

There are several products for which you can specify additional output options. If one of these products is on the screen, the **Output Options** tool icon is enabled.

The **Output Options** tool allows you to adjust the display of these products for your application.

You can then select **File > Save** for the **VVP**, **WIND/FCAST**, and **NDOP** options to save the configuration for your window is saved (it does not affect other windows).

> If you are using **WIND**, **FCAST**, or **NDOP** products as overlays on other products, it is sometimes convenient to display (for example) the **WIND**, product and the options tool for **WIND**. Then leave the tool up so you can use it if needed when you look at **WIND** overlaid on another product.

6.17.1 CAPPI Height Selection Tool

For **CAPPI** products with multiple levels (3D **CAPPI** products), you can use this tool to change heights.

You can also use the **UP**/**DOWN** arrow keys (your cursor must be in the display portion of the window to use the arrow keys).

More Information

‣ CAPPI: Constant Altitude Plan Position Indicator (page 39)
‣ NDOP Output Options (page 256)

6.17.2 THICK Output Options

Figure 90   **THICK** Output Tool

The **THICK** product shows the difference in height of the lowest and highest instance of any reflectivity contour.

The **THICK** product may also show the average reflectivity value within the same layer. This tool allows you the change the data types is displayed in the **THICK** product.

6.17.3 NDOP Output Options

NDOP product output options are a combination of the **WIND** options and the **CAPPI** height selection.

Because this tool works with **WIND**, **FCAST**, **NDOP**, and **CAPPI**, you may wish to leave it on the screen.

More Information

‣ NDOP : Multiple Doppler (page 174)
‣ CAPPI Height Selection Tool (page 255)
‣ WIND and FCAST Output Options (page 264)

6.17.4 VAD Output Options

During display, you can select one or all of the elevation angles in the **VAD** product.

The following figure shows an example of a single angle being displayed. This data has not been unfolded.

Figure 91   Example of Velocity Azimuth Display

Figure 92   **VAD** Output Tool

All Heights/1 Height

> The IRIS **VAD** products consist of a plot of velocity for all azimuths. You can choose to display just one plot of one height, or many plots showing all the heights.

Height

> If you have selected **1 height**, use this slide bar, or the cursor keys to move up and down through the list of heights.

Plot Max

Configure here the maximum velocity shown on the plots.

More Information

‣ VAD: Velocity Azimuth Display (page 97)

6.17.5 VVP Output Options: Time-Height

You can use the **Time-Height Output Options** to show a height against time cross-section of **VVP** parameters with the specified color background and overlay plot.
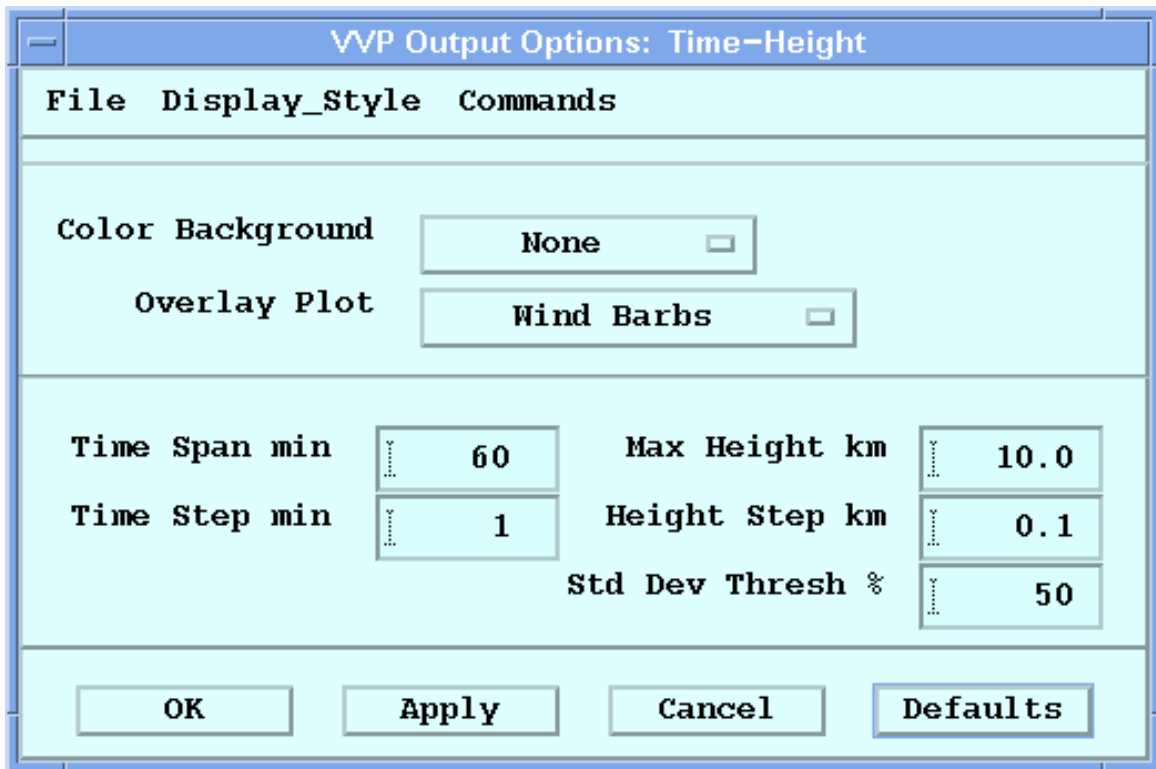
The **VVP** product shows vertical profiles of average Doppler wind properties over the radar such as wind speed, wind direction and divergence. The **Output Options** menu lets you choose to display wind data in one of the following ways:

- Choose **Display_Style > Graphs** to display a graph of **VVP** parameters against height at a single time. You can select 1, 2, or 3 graphs.
- The **WIND** product plots the horizontal wind speed and direction.
  For this product type, you can choose whether to display all wind measurements or only those winds that deviate from the mean.
  You can choose either wind barbs or wind strings to represent wind speed and direction, and you can specify a coverage threshold percentage.

1. Check that **VVP** is the current product type.

2. Select **Options**.
   The **Time-Height Output Options** menu is shown.

3. In **Color Background**, choose the information that you want displayed in the background.

```
Divergence
Deformation
Vert Velocity
Reflectivity
None
```

4. In **Overlay Plot**, Choose the type of overlay plot.

```
Wind Barbs
Wind Strings
Axis of Dilatation
None
```

   - Wind barbs use the standard meteorological convention of 1/2 barb for 5 knots, 1 barb for 10 knots and triangles for 50 knots.
   - Wind strings point away from the wind direction, with the length of the string indicating the relative wind speed.
     You can set the maximum length of the wind strings using the **Display Style > Graphs Scale Limit** for wind speed.
   - **Axis of Dilatation** lines show the angle of "stretching".
     Fronts tend to form parallel to this angle. The length is proportional to the deformation which indicates the strength of the stretching.
     You can set the maximum length of the dilatation using the **Display Style > Graphs Scale Limit** for deformation.
   - Select **None** if you do not want to display an overlay plot.

5. In **Time Span min**, set the duration of the time (horizontal) axis in minutes.

6. In **Time Step min**, set the minimum time step for plotting data.
   This is useful to avoid collisions of the wind barbs when the data are closely spaced. Type **0.0** to see all the data.
   For example if there is a **VVP** profile every 5 minutes and this field is set to 10, then every other profile is plotted.

7. In **Max Height km**, type the maximum height (vertical scale) that you want to display. Note that the height of the **VVP** profile may be less than the height that you display, in which case a blank background color is shown.

8. To avoid collision of the wind barbs in height, you can specify **Height Step km**.
   Type `0.0` to see all the data.

9. In some cases, the wind estimates are noisy because of weak signals or poor coverage.
   If needed, use **Std Dev Thresh %** to threshold poor estimates so that they are not
   displayed.
   Specify a number between 0 (no thresholding) and 100 (everything thresholded away).
   Typically a value of 80 is used.

10. To test your setting, select **Apply**.
    The velocity point is displayed if the following is true:

$$Stan \, dardDeviation < \; -In\left(\frac{threshold}{100}\right) \times (1.0 + Velocity/10.0)$$

More Information

‣ Producing VVP Product Output (page 104)
‣ VVP Output Options: Line Graphs (page 260)
‣ Color Scale Tool (page 217)
‣ WIND and FCAST Output Options (page 264)

6.17.6 VVP Output Options: Line Graphs

VVP products can be displayed in graphical format of **VVP** wind profile parameters against
height. 1, 2, or 3 graphs may be plotted side-by-side.

1. Use the **Product Selection** tool to select a **VVP** product.

2. Select the **Options** icon to show the **Display Options** menu.
   Make sure the **Product Type** is **VVP**.

3. Select **Options** to show the **VVP Output Options** menu.
   Either the **Graphs** or the **Time Height** submenu is shown.

4. If **Time Height** is shown, select **Display_Style > Graphs**.

5. Make your changes by enabling the plot parameters you want to use and typing relevant values in the **Scale Limits** column:

Wind Speed

Select to chart the wind speed in the graph.

Type a value to specify the maximum wind speed you want to plot.

Wind Direction

Select to chart the wind direction. You cannot edit the **Scale Limits** column for this field.

Divergence

Select to chart the wind divergence. Type a value to specify the maximum divergence you want to plot.

Vertical Velocity

Select to chart the vertical velocity. Type a value to specify the maximum vertical velocity you want to plot.

Deformation

Select to chart the horizontal deformation. Type a value to specify the maximum deformation you want to plot.

Axis of Dilation

Select to chart the axis of dilatation. You cannot edit the **Scale Limits** column for this field.

Reflectivity

Select to chart the reflectivity. Type a value to specify the maximum reflectivity you want to chart.

Max Height in km

Type the maximum height for the data to include in the plot.

Std Dev Thresh %

Enter the standard deviation threshold percent for the data to include in the plot. This is used to threshold poor estimates. The scale goes from 0 (no thresholding) to 100 (total thresholding). 80 is a typical value.

> ℹ️ Alternatively, select **Defaults** to use a reasonable set of default values as a starting point.

6. Select **Apply**.

Your changes are shown in the current **VVP** display as a graph.

7. Select **Display Options > File > Save**.

More Information

6.17.7 WARN Output Options
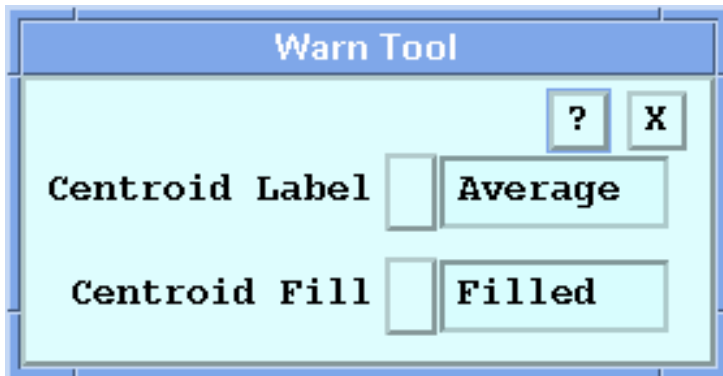
Figure 93   **WARN** Output Options

IRIS **WARN** centroids are displayed as an ellipse with a name inside or as a graphical icon rotated to the direction of motion of the centroid.

Centroid Label

The **Centroid Label** options are for a second label line on ellipses or for a label to the right of the icon.

The **WARN** centroid has 4 label choices. The data displayed is from the first of the 3 possible inputs to the warning. The units are the same as in a standard IRIS data display, except that wind shear is displayed as knots loss/gain, and height is displayed in kilofeet.

None
No additional label.

Maximum
The centroid is labeled with the maximum value of the data throughout the centroid.

Average
The centroid is labeled with the average value of the data throughout the centroid.

Speed
The centroid is labeled with the speed of a target. This only applies to systems that have the **DWELL** product configured for target detection.

Centroid Fill

The **Centroid Fill** options apply to ellipse-style displays when the **WARN** product is overlaid on other products.

You can choose the style for the fill effect of the centroid so that it can be clearly identified on your display. The options are:

Open
The centroid area is not filled, it is outlined.

Hatched
The centroid area is filled with a slanted pattern of lines.

Filled
The centroid area is filled a solid color.

Operational Use

Configure your **Warn Tool** options and then, in **Display Options**, select a product to overlay.

Select **File > Save** to save the **Warn Tool** options.

6.17.8 WIND and FCAST Output Options
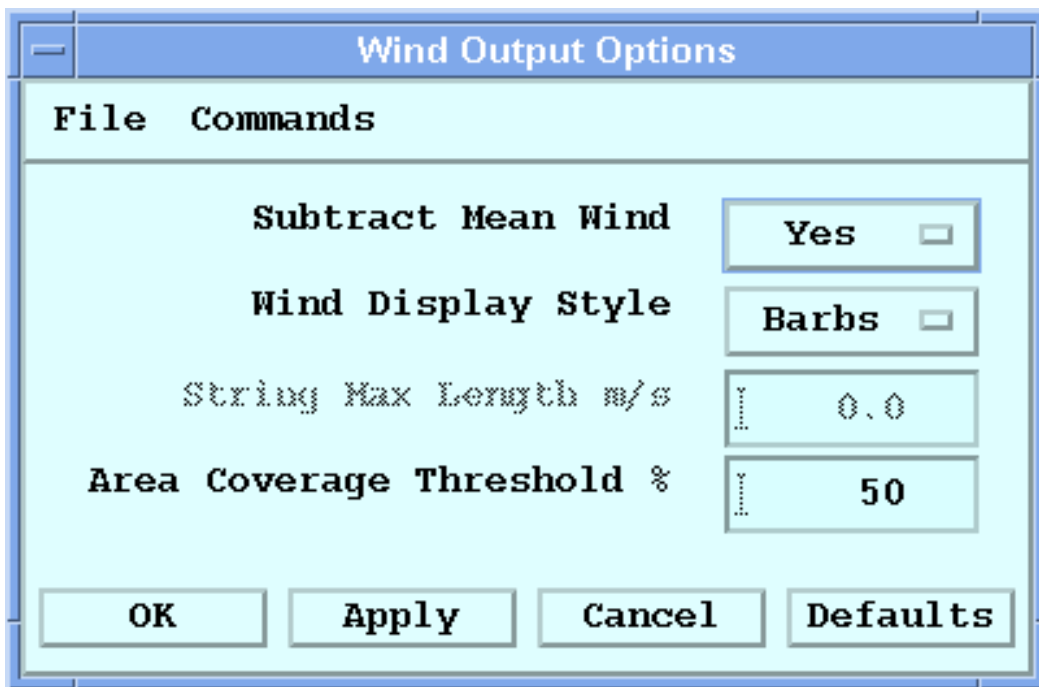
1. Make sure **WIND** or **FCAST** is the current product type.

2. Select **Options**.
   The **Wind Output Options** menu appears with options for plotting the horizontal wind speed and direction.

3. Define the **WIND** display options.

Subtract Mean Wind

Select **Yes** for IRIS to subtract the mean wind vector from all wind vectors and display the perturbation wind vectors.

Select **No** to display the full wind vector.

Wind Display Style

Display the **WIND** product vectors as wind barbs or wind strings.

The **WIND** vectors are estimates of horizontal winds using a uniform wind algorithm in sectors around the radar.

Wind barbs point toward the wind direction, with hatch marks denoting the wind speed following the standard meteorological convention of 1/2 barb for 5 knots, full barb for 10 knots, and triangles for 50 knots.

Wind strings point away from the wind direction, with the length of the string indicating the relative wind speed (the longer the string, the higher the speed).

String Max Length

If you choose **Strings** as the wind display style, enter the maximum string length in meters/second.

Any data point that exceeds this length is displayed at the maximum length.

Area Coverage Threshold %

Add threshold points for which the aerial coverage of radial velocity data is less than the selected value. A value of 50 % is typical.

More Information

‣ NDOP Output Options (page 256)
‣ VVP Output Options: Time-Height (page 258)
‣ FCAST: Forecast (page 40)
‣ WIND: Wind Speed and Direction (page 113)

6.18 Printing and Exporting Displays

> **ℹ** For information on automatic printing and exporting see 7.4 Sending a Product to a Device (page 274).

1. To export a display, select **File > Export**
   The **Export Tool** appears with settings for outputting to a disk file on your local computer or to a remote computer in GIF, TIF, BMP, or Post Script formats.
   The formats are controlled in **Setup > Output**. See *IRIS and RDA Utilities Guide*.
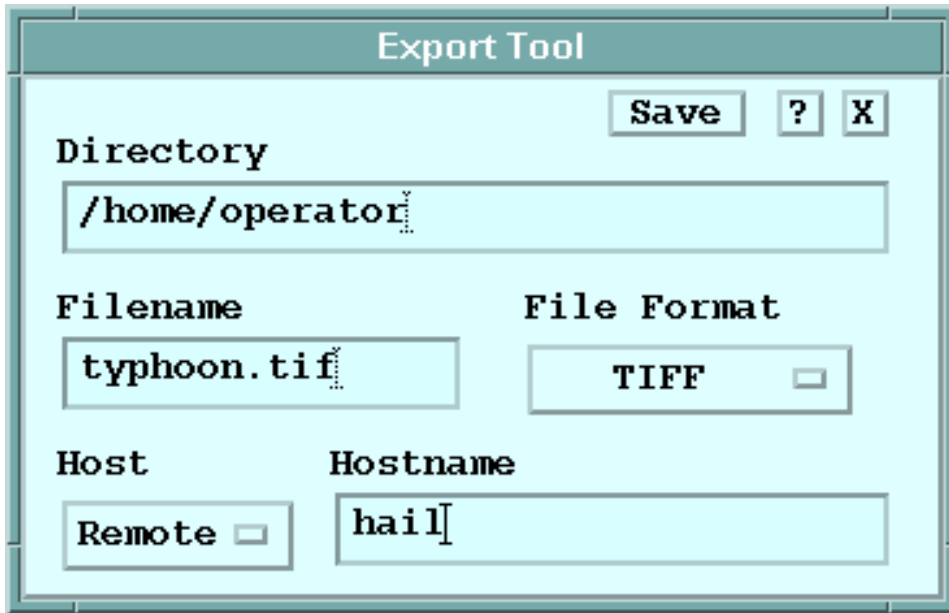
2. To print a display, select **File > Print**
   For information on setting-up postscript printing, see *IRIS Radar User Guide*.

6.19 QLW Keyboard Commands

> **💡** To use the keyboard commands, place your cursor in the image area of the Quick Look Window (QLW).

Keystroke	Action
↔ LEFT/RIGHT Arrow	Forward/Back in single-step time sequence
SPACE BAR	Toggle loop on/off
↕ UP/DOWN Arrows	Up/Down height step for **CAPPI** or **NDOP**

7. Requesting Product Output

7.1 Overview of Requesting Product Output

The product files created by the product generator are stored on disk. From there, they can be sent to displays, printers, storage media, or other computers.

The output process adds any requested overlays, then formats the product for the selected output device. The **Product Output** menu directs the output to the required destination.

You can use the **Product Output** menu to browse the available product files and select some to send to your workstation. You can also request that all future versions of a particular product be sent automatically to your workstation or selected device. For example, you might want to see a **PPI** product every time a particular volume scan task runs. In addition, an operator might send a product to a particular destination, such as a remote display, printer or DVD.

7.2 Product Output Devices and Files

The **Product Output** menu can route output to many devices:

Printer
Postscript printers configured on you system.

Window
A color window on the local workstation. The display is generated by the IRIS host. This window can be exported to appear on a remove workstation.

Network
Sends product files to another system over the network or to a disk file on the system. The remote system can receive a socket notification of the file. The file format can be converted to any other format using "output pipes".

Archive
A tape drive or MO disk can be used to store data.

You configure the list of output devices in the **Setup** utility. See *IRIS and RDA Utilities Guide*.

For a list of output pipes supplied with the system see *IRIS and RDA Utilities Guide*.

Cartesian products can be displayed in a window using up to 32 data colors, with additional colors available of overlay and annotation. Any product resolution can be displayed. The standard IRIS display sizes include the following resolutions:

- **Medium**. For 480 x 480 pixel displays.
- **High**. For 720 x 720 pixel displays.
- **X-High**. For 940 x 940 pixel displays.

7.3 Accessing the Product Output Menu
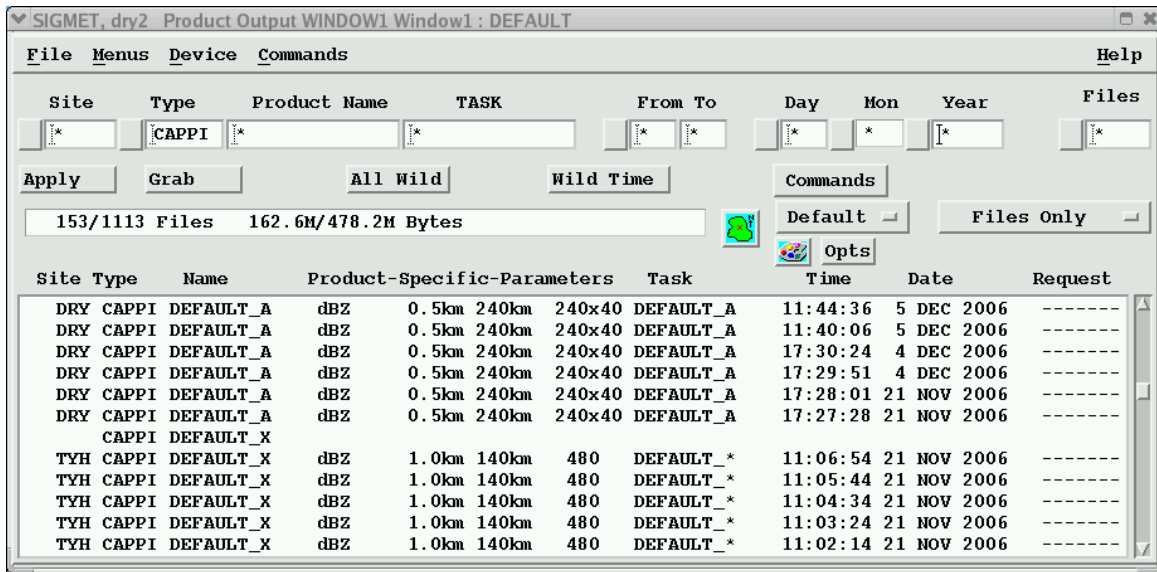
Filter Menu
>   Lets you choose the files that appear in the product list.
>   See 7.3.1 Product Output Filters (page 268)

Product List
>   Displays the list of products that you have selected.
>   See 7.3.2 Product Output: Product List (page 271)

▶   1. Choose **Menus > Product Output** from the IRIS menu bar or from any of the IRIS
        menus.

7.3.1 Product Output Filters

Figure 94   Product Output Filters

When you first enter the **Product Output** menu, all possible products are listed at the
bottom of the menu. This list can be quite large, making it difficult to find the products you
need. Use the **Filter** section to narrow the list down to only those products that are of
interest.

Site

Enter a site ID to select data from only one radar site or enter the wildcard character to select data from all sites.

Select **Site** to show a list of known sites. The contents of this list depend on the sites that are configured in the **Setup** utility.

Type

Enter a product type in this field to select products of only one type, or enter the wildcard character to select products of all types.

You can also select from a list of types with data files.

Product Name

Enter the name of a product to narrow the list down to products having that name. You can include wildcard characters in the product name. A question mark (?) matches any single character; an asterisk (*) matches any sequence of zero or more characters.

TASK

Enter the name of a task to filter for products generated by a specific task.

You can include wildcard characters in the task name.

For example, entering a product name of `*PPI*_?` selects all hybrid tasks that have `PPI` somewhere in their names.

Scan

Enter a scan type in this field to select products of only one type, enter the wildcard character to select products of all types, or select from a list of types.

From, **To**

Enter a range of hours to filter products by time of day.

You can also select from a list of ranges.

Day, **Month**, **Year**

Enter a day, month. or year to filter products by date.

You can also select from a list of values.

Files

Control how many files to include in the list.

Max

Control how many products to include in the list. "*" means apply no limit. All IRIS scrolled lists are capped at 1000.

Apply

Select **Apply** to update the product list, based on your selection criteria.

Grab

If you have selected a product, select **Grab** to insert the product's site, type, product name, task, and date information in the filter fields.

All Wild

Select **All Wild** to return all the fields to the wildcard character.

Wild Time

Select **Wild Time** to change only the hours, month, day, and year fields to the wildcard character.

Commands

Select one of the following operations to perform on product files selected by the filter menu. These commands apply to products, but not to headers:

- **Keep**
Prevents the Watchdog process from deleting the selected products.
- **Delete**
Marks the selected files for deletion by the Watchdog process.
- **Reingest**
Marks the selected files to be reingested.
- **Remove Tag**
Unflags any selected products that have been tagged.
- **Send**
Sends all selected products to the output device.
- **Cancel Send**
Cancels any send requests for the selected products.

Display Options

Select the this icon to specify the output options to be used when your product is sent to a display device.

Color Scale Configurations

Select this icon to specify the color options to be used when your product is sent to a display device. See 6.7 Color Scale Tool (page 217).

Product Specific Output Options

Opts

This button is grayed out unless the **CAPPI**, **FCAST**, **NDOP**, **THICK**, **VVP**, or **WIND** product type is selected.

- 3.5 CAPPI: Constant Altitude Plan Position Indicator (page 39)
- 3.6 FCAST: Forecast (page 40)
- 5.5 NDOP : Multiple Doppler (page 174)
- 3.19 THICK: Echo Thickness Product (page 89)
- 3.24 VVP: Velocity Volume Processing (page 101)
- 3.26 WIND: Wind Speed and Direction (page 113)

Show Headers

Show header lines with one of the following options:

- **Files Only**
Show files in the inventory with a header for each file shown. This also shows **Headers in Use**.
- **Headers Only**
Show header lines only.

- **Files + Headers**
  Show files and all headers.
  The only additional headers seen in **Files + Headers** are for products on disk which you may have recently deleted.
- **Headers in Use**
  Show headers which have automatic send requests configured.
  This shows headers with automatic send requests for any output, not just the output device currently selected.

> You may want to set up your product output assignments based solely on product headers. Select **Headers Only** to do this.

7.3.2 Product Output: Product List

```
Site Type     Name       Product-Specific-Parameters     Task          Time      Date        Request
  DRY CAPPI DEFAULT_A     dBZ    0.5km 240km   240x40 DEFAULT_A    11:44:36  5 DEC 2006    -------
  DRY CAPPI DEFAULT_A     dBZ    0.5km 240km   240x40 DEFAULT_A    11:40:06  5 DEC 2006    -------
  DRY CAPPI DEFAULT_A     dBZ    0.5km 240km   240x40 DEFAULT_A    17:30:24  4 DEC 2006    -------
  DRY CAPPI DEFAULT_A     dBZ    0.5km 240km   240x40 DEFAULT_A    17:29:51  4 DEC 2006    -------
  DRY CAPPI DEFAULT_A     dBZ    0.5km 240km   240x40 DEFAULT_A    17:28:01 21 NOV 2006    -------
  DRY CAPPI DEFAULT_A     dBZ    0.5km 240km   240x40 DEFAULT_A    17:27:28 21 NOV 2006    -------
      CAPPI DEFAULT_X
  TYH CAPPI DEFAULT_X     dBZ    1.0km 140km      480  DEFAULT_*    11:06:54 21 NOV 2006    -------
  TYH CAPPI DEFAULT_X     dBZ    1.0km 140km      480  DEFAULT_*    11:05:44 21 NOV 2006    -------
  TYH CAPPI DEFAULT_X     dBZ    1.0km 140km      480  DEFAULT_*    11:04:34 21 NOV 2006    -------
  TYH CAPPI DEFAULT_X     dBZ    1.0km 140km      480  DEFAULT_*    11:03:24 21 NOV 2006    -------
  TYH CAPPI DEFAULT_X     dBZ    1.0km 140km      480  DEFAULT_*    11:02:14 21 NOV 2006    -------
```

The product list at the bottom of the **Product Output** menu lists each product name on a header line, followed by the versions of the product that are currently on disk. Each version is identified by the time the data were collected, which is the start time of the associated task.

Site and Type

The **Site** field displays the ID of the radar site where data for the product was collected.

The **Type** field shows the type of the product, such as **PPI** or **RHI**.

> You can right-click the fields to manage the files. See .

Name

The product name for the configuration of this product type.

Product-Specific Parameters

The standard product generation parameters shown in the following table are taken from the product configuration and displayed in **Product-Specific Parameters**.

Table 21   Standard Product Generation Parameters

Product Type	Parameter 1	Parameter 2	Parameter 3
BASE	...	dBZ Threshold	Max Range
BEAM	Azimuth Limits	Elevation Limits	...
CAPPI	Data Parameter	Height	Max Range
CATCH	Hours	...	Max Range
FCAST	Speed & Direction	Data Threshold	Max Range
GAGE	Minutes avg	Rain/Disdro	...
IMAGE	...	...	Max Range
MAX	Layer Bottom	Layer Top	Max Range
NDOP	List of Sites	...	Max Range
PPI	Data Parameter	Elevation	Max Range
RAIN1	Input Product	Minutes avg	Max Range
RAINN	Hours	...	Max Range
RAW	Data Parameters	All/One	Max Range
RHI	Data Parameter	Azimuth	Max Range
RTI	Data Parameter	Start Az and El	Max Range
SHEAR	R, A, or E flags	Elevation	Max Range
SLINE	Elevation(s)	...	Max Range
SRI	Melting level	Profile/Convection flags	Max Range
TDWR	CF Wind Speed	CF Wind Dir	CF Gust Speed
TOPS	...	dBZ Threshold	Max Range
TRACK	Inputs Product	Area Threshold	Data Threshold
USER	Data Parameter	...	Max Range
VIL	Min Height	Max Height	Max Range
VVP	Min & Max Height	Number of Steps	Bin Quota
WARN	Symbol & Area	Number of Inputs	Number of hits
WIND	Min & Max Height	...	Max Range
XSECT	Data Parameter	Ctr Point R, AZ	Orientation Angle

Resolution

The product-specific parameters for Cartesian products include the horizontal resolution of the product in pixels. For 3-D products, the vertical resolution is also listed, for example, 480×10.

The product output process always generates a display even if the output resolution does not match your display capabilities. It does this by skipping or duplicating pixels as required. For example, if you send a low resolution display to a medium resolution display device, such as a printer, the Product Output process duplicates each pixel to zoom the image on the printer to full size.

TASK

The name of the associated task.

Time

The date and time (local time) when the product data was collected.

Request

Use the **Request** field to request product output. Position the mouse cursor over an entry in the product list and select **Send**.

The menu shown depends on the entry you select in the list.

If you select a header line from the list, you can choose to send all or only the next occurrence of the product to the output device. You can also select particular sites that you want to receive the output. Each site you select is marked with an asterisk.

After you issue the **Send** request, the **Request** field displays a message indicating the status of the request.

7.4 Sending a Product to a Device

You can send one or more existing products to a display device, or one or more future versions of a product to a display device at one or more sites.

▶ 1. To send existing products:

a. Select **Device**.

b. Select the device to which you want to send the product.

c. Select the products you want to send by doing one of the following:
   - Left-click the mouse over a list item.
     Press **CTRL**+click to select more than one entry in the list.
   - Click and drag the mouse to select a group of entries in the list.
     Press **CTRL**+click and drag the mouse to select multiple groups of entries.
   - Use the **Filter** menu to produce a list of the products that you want to send.

d. Choose **Send** from one of the following menus:
   - If you selected the entries with the mouse, position the cursor over the product list and right-click the mouse.
   - If you selected the entries from the **Filter** menu, select **Commands**.

e. Select **Send**:
   The product is marked to be sent to the selected device and is waits for the output formatter to send it.
   The sending status is shown with the following indicators:

   Sending
   > The output formatter is working on your request.

   Done
   > The product has been sent to the selected display device.

   > There is currently no request.

   Error
   > An error occurred. Check the **Radar Status** menu.

> 💡 If the selection contains both headers and products, the command applies only to the products.
> To perform an operation on headers, you must select only headers.

2. To send one or more future versions of a product:

   a. Select the header for the type of product you want to send.

   b. Optionally, select any number of sites that are known to your system to receive all or the next version of the product.

   c. Select **All** or **Next**.

   A send request is issued automatically when a new version of the product is added to the disk.
   A send status message appears in the **Request** column.

3. To cancel a send request, place the cursor over the **Request** column and select **Cancel**.
   If you select a single product, you cancel only the send request for that product.
   If select the header line, you cancel the **All** or **Next** request.

7.5 Flagging a Product

Normally, the Watchdog process removes product files as the disk space fills by deleting the oldest product files.

You can also signal the Watchdog to keep or delete product files.

▶ 1. In the first column of the product list, left-click the mouse over a list entry.
   To select more than one entry in the list, press **CTRL**+click.

2. Use the **Filter** menu to narrow the list down to the entries you want to flag.

3. Flag the entries with one of the following options:
   Keep
   > Prevents the Watchdog from deleting a file.
   > This is useful if you have files that you want to save for several days.
   > Note that the best way to make a record of a file is to send it to tape.

   Delete
   > Signals the Watchdog to delete a file immediately.
   > This is useful if you want to restore a RAW file from tape so that it is automatically reingested.
   > This cannot be done if the RAW product file already exists on disk. It must be deleted first.

   Clutter
   > For **RAIN1** products, this flags the file that serves as the clutter map for the product.
   > If more than one file of the same product name is tagged, the most recent version is used. The Watchdog does not delete clutter maps.

   Remove Tag
   > Removes a flag.

8. Performing Archive Operations

8.1 Overview of Archive and Retrieve Operations

IRIS can archive, retrieve, and play radar data. If a RAW product (polar coordinate data) is retrieved from archive, IRIS can generate products from the retrieved data as if the data were collected in real time. Picture products can be retrieved from archive and then displayed. Supported devices include:

- Tape such as DAT
  Recording to tape is convenient, but retrieval can be slow if the file that you want is at the end.
- Optical disk drives
  Optical disks, such as re-writable magneto optical drives, are more expensive than tape, but last longer and are much more convenient to use because of their random access capability.
- Large Disk Archives (LDA)
  LDAs allow you to take advantage of low-cost mass storage systems.
- DVD+RW systems (for example, SONY DRU–500AX (internal), SONY DRX–500ULX (external on Firewire or USB 2.0).
  DVDs offer low cost for both the drive and the media while providing approximately 5 GB of storage. Supported on Linux systems.

8.2 Using the Archive Menu

The **Archive** menu controls the host archive devices.

wind Archive Menu TAPE1 DAT									

File   Menus   Drive   Commands   **Help**

Site   Type   Product Name   TASK   From To   Day   Mon   Year   Files

[   ] [   ] [   ] [   ]   [ ][ ] [ ] [ ] [ ]   [100]

Apply   **Grab**   **All Wild**   **Wild Time**   **Commands** **Start List at** [1]

36 /36 Files 9.7M/9.7M Bytes

Commands [Mount ⌐] [Done]   **Status** [Idle @ 1]   **Write** [On]

Archive ID [SIG_96_11_27_E1W]

File	Site	Type	Product	Products-Specific-Parameters			Task	Time	Date
d	1	HAI	CAPPI	Z_020_240 Z	2.0km 240km	480 P	PPIVOL	17:11:13	12 Jul 1997
d	1	HAI	PPI	Z_020_240 Z	2.0deg 240km	480 P	PPIVOL	17:11:13	12 Jul 1997
d	1	HAI	CAPPI	Z_030_240 Z	3.0km 240km	480 P	PPIVOL	17:11:13	12 Jul 1997
r	1	HAI	CAPPI	V_020_240 V	2.0km 240km	480 P	PPIVOL	17:11:13	12 Jul 1997
r	1	HAI	CAPPI	V_030_240 V	3.0km 240km	480 P	PPIVOL	17:11:13	12 Jul 1997
d	1	HAI	RHI	Z_XXX_120 Z	90.0deg 120km	600 D	RHI	16:10:13	12 Jul 1997
d	1	HAI	RHI	V_XXX_120 V	90.0deg 120km	600 D	RHI	16:10:13	12 Jul 1997
	1	HAI	RAW	AERIAL	Z V W 240km	One	AERIAL	16:05:00	12 Jul 1997

Figure 95   **Archive Menu**

Filter Menu
> You can use the standard filter menu for selecting the products that appear in the Archive Log area.

Archive Control Area
> The middle of the menu provides commands to control archive operations such as mount or record, and displays status information for the archive device functions.

Archive Log Area
> The bottom of the menu shows what product files, if any, are recorded on the archive that is currently mounted.

1. Make sure the archive device is configured in **Setup > Archive**.

2. In IRIS, select **Menus > Archive Menu** and, for example, do the following:
   • Mount an IRIS tape or disk (archive media) and start the recording process.
   • Retrieve product files and store them on the system disk.

3. All operations are directed to a selected archive device. To select an archive device (tape or optical disk), select **Archive > Drive**.

4. Select a drive from the list.
   The choices vary depending on the drives configured for your system. If no archive drives are available, the list contains the entry `UNKNOWN`.

5. To direct output to the archive device, make send requests in **Menus > Product Output**.

8.2.1 Archive Control Area

Figure 96   Archive Control Area

The archive control area lets you select the archive device commands and check status information.

Commands

Select **Commands** to show the following tape commands:

- **Mount** a disk or tape.
- **Unmount** a disk or tape.
- **Stop** the current operation.
- **Record** a product file to archive.
- **Retrieve** a product file from the archive and copy it to disk.
- **Inventory** the product files on the archive.

Status

Status shows the current status of the archive drive and should correspond to the command request. The status entries are:

- **Rewind** (for tape)
  The tape is rewinding, for example:
  - The user issues the **Unmount** command.
  - IRIS is positioning the tape for a **Retrieve** operation.
  - IRIS has filled a tape during a **Record** operation.
- **Writing**
  The tape or disk is being recorded.
- **Scanning**
  The tape is being mounted or an inventory is being created. **Scanning** is also displayed while the tape is positioning for a **Retrieve** or **Record** operation.
- **Reading**
  The drive is in the process of retrieving a file.
- **Idle**
  The drive is not doing anything, that is, the user did not issue a command, or the last command was completed.

- **Off Line**
  The specified drive is not online. No archive operation can be performed until this problem is corrected. You may need to insert the medium or push a load button. See the drive manual.
- **Unavailable**
  The drive is being used by someone else, it is broken or not installed.
- **Released**
  IRIS is not currently attached to the drive. The drive is free to be used by IRIS or other processes.
- **Full**
  The medium is full, an EOT or disk full was encountered during a **Record** operation. To continue recording, the medium must be unmounted, a new medium mounted and **Record** must be invoked again.
- **BOT** (tape)
  Beginning of tape, a new tape has been inserted into the drive and is ready to be mounted. This status also appears after a **Rewind** is done.
- **EOT** (tape)
  End of Tape, IRIS has encountered the end of tape. Usually IRIS rewinds the tape automatically, so the message is replaced by **Rewind**.

Up to 3 status indicators can appear together in the **Status** field. For example, if a **Record** process fills the tape, the tape is rewound and sits idle until the operator changes tapes or issues another command, the **Status** field indicates the following:

```
Idle BOT Full
```

If the tape is being scanned to create an inventory, the status field displays the following:

```
Scanning
```

When a record process currently has nothing to do, the status displays:.

```
Idle
```

If new product files tagged for archive (in the **Product Output** menu) become available, they are recorded automatically.

When the **Mount** command is completed and the tape is at the beginning awaiting commands, the **Status** field displays:

```
Idle BOT
```

Archive ID

Identifies the unique name assigned to the tape or disk when it was initialized.

The first 3 characters are the IRIS site ID for the site where the medium was initialized, followed by the date on which the medium was initialized. The final 3 characters are a unique random code.

Write

Write shows whether the write protection for the tape or disk is set.

- **On** means that IRIS can write on the medium.
- **Off** means that IRIS cannot write on the medium.

The message appears after a **Mount** command is issued.

Filter Menu Commands

Site

Enter a site ID to select data from only one radar site or enter the wildcard character to select data from all sites.

Select **Site** to show a list of known sites. The contents of this list depend on the sites that are configured in the **Setup** utility.

Type

Enter a product type in this field to select products of only one type, or enter the wildcard character to select products of all types.

You can also select from a list of types with data files.

Product Name

Enter the name of a product to narrow the list down to products having that name.

You can include wildcard characters in the product name. A question mark (?) matches any single character; an asterisk (*) matches any sequence of zero or more characters.

TASK

Enter the name of a task to filter for products generated by a specific task.

You can include wildcard characters in the task name.

For example, entering a product name of `*PPI*_?` selects all hybrid tasks that have `PPI` somewhere in their names.

Scan

Enter a scan type in this field to select products of only one type, enter the wildcard character to select products of all types, or select from a list of types.

From, **To**

Enter a range of hours to filter products by time of day.

You can also select from a list of ranges.

Day, **Month**, **Year**

Enter a day, month. or year to filter products by date.

You can also select from a list of values.

Files

Control how many files to include in the list.

Max

Control how many products to include in the list. "*" means apply no limit. All IRIS scrolled lists are capped at 1000.

Apply

Select **Apply** to update the product list, based on your selection criteria.

Grab

If you have selected a product, select **Grab** to insert the product's site, type, product name, task, and date information in the filter fields.

All Wild

Select **All Wild** to return all the fields to the wildcard character.

Wild Time

Select **Wild Time** to change only the hours, month, day, and year fields to the wildcard character.

Select **Commands** to show operations you can perform on the files from the **Filter** menu:

- **Retrieve** marks all selected files for retrieval.
- **Cancel** unmarks all selected files.

Filter Menu Start List At

The Archive log listing can contain several hundred files. Use the **Start List At** option to start the list at a selected file.

You can select a starting number from a menu or enter a number directly into the field. This is convenient for tapes that are in chronological order.

8.2.2 Archive Log Area

File	Site	Type	Product	Product-Specific-Parameters	Task	Time	Date
d 1	HAI	CAPPI	Z_020_240 Z	2.0km 240km 480 P	PPIVOL	17:11:13	12 Jul 1997
d 1	HAI	PPI	Z_020_240 Z	2.0deg 240km 480 P	PPIVOL	17:11:13	12 Jul 1997
d 1	HAI	CAPPI	Z_030_240 Z	3.0km 240km 480 P	PPIVOL	17:11:13	12 Jul 1997
r 1	HAI	CAPPI	V_020_240 V	2.0km 240km 480 P	PPIVOL	17:11:13	12 Jul 1997
r 1	HAI	CAPPI	V_030_240 V	3.0km 240km 480 P	PPIVOL	17:11:13	12 Jul 1997
d 1	HAI	RHI	Z_XXX_120 Z	90.0deg 120km 600 D	RHI	16:10:13	12 Jul 1997
d 1	HAI	RHI	V_XXX_120 V	90.0deg 120km 600 D	RHI	16:10:13	12 Jul 1997
1	HAI	RAW	AERIAL	Z V W 240km One	AERIAL	16:05:00	12 Jul 1997

Figure 97   Archive Log Area

The archive log area shows a list of files that are on the archive. This list is used for specifying files for a retrieve operation or verifying the progress of a record operation.

For tapes, the order of the list is the same as the order of the files on the tape. Usually, they are in chronological order.

For optical disks, the files appear in the same order as in the **Product Output** menu.

Archive logs are maintained on disk as ASCII text files.

When a tape is mounted, IRIS checks to see if there is a corresponding tape log. If the tape is a valid IRIS tape, the list is displayed.

In the case of optical disk, a new log is generated when a disk is mounted.

File

This column contains status information about the file. The first column displays the sequential tape file number (1, 2, ...) and the second column contains one the following symbols:

- r
  The file has been tagged for a retrieve operation.
- d
  The file is on disk as well as on tape. Retrieve is not permitted.
- "blank"
  The file is on tape only and has not been tagged for retrieve.

Right-click **File** to select one of the following options:

- **Retrieve** marks a file for retrieval.
- **Cancel** unmarks the file.

Site

Shows the ID of the radar site where data for the product was collected.

Type

Shows the type of the product, such as **PPI** or **RHI**.

Product

This display-only field gives the product name for the most recent generation of the product type.

Product-Specific-Parameters

Shows additional information about the product. See Table 21 (page 272).

TASK

This display-only field shows the name of the associated task.

Time

This display-only field shows the date and time (local time) when data for the product was collected.

8.3 Initializing a Tape or Disk for Recording

IRIS can mount only valid IRIS tapes or disks that have been initialized by a separate utility program. If you try to mount a non-initialized medium, the **Status** field shows:

```
NonIris
```

Use the medium initialization utility to do the following:

- Initialize a blank tape or disk.
- Overwrite a tape or disk that contains non-IRIS data.
- Overwrite an IRIS tape or disk.

> ⚠️ **CAUTION!** If you are overwriting an IRIS or other type of medium, make sure that you are not destroying valuable data or software (backup or delivery tapes).
> Only operators are permitted to use the tape initialization utility.

> ⚠️ **CAUTION!** Do not initialize an existing tape or disk if you want to record additional data or retrieve IRIS data. Initializing an existing IRIS tape or disk destroys the data already on the medium.

▶ 1. Make sure the drive is free. Any IRIS tape or disk must be unmounted and removed.

2. Insert the medium into the drive and set the drive on line. Be sure that the medium is write-enabled.

3. If IRIS is running, check the **Archive** menu and, if necessary unmount the media that you want to initialize.

4. From the operating system prompt, invoke the **init_iris** command appropriate for your device.
   Here are some examples:

```
$ ini t_iris_tape  (for drive unit 1 tape)
$ init_iris_mo     (for drive unit 1 disk)
$ init_iris_mo -u2 (specifies the 2nd drive unit)
$ init_iris_dvd -h (to get help on usage)
$ init_iris_lda -h (to get help on usage)
```

Be sure to type the underscore (_) characters as shown. The command name is intentionally long so you cannot type it accidentally.
The "−u2" specifies the second drive if there is more than one.
The "−h" argument prints the usage options.
The unit is checked and you are prompted to confirm the initialization. The drive goes through several initialization steps that are displayed on the screen. If you have forgotten to write enable the tape, you get a write protect error when IRIS attempts to write the header information. Write-enable the tape and retype the command.
At the end of the initialization, IRIS displays a unique tape number that is used in the future to refer to the disk inventories maintained for the tape.

> ⚠️ **CAUTION!** Write the name of the medium that IRIS assigns on the tape or disk label. This lets others know that the tape or disk has been initialized.

The tape name is based on the first 3 characters of the Site ID, the year, month, day and a random number to make the tape unique.

8.4 Mounting a Tape or Disk

Use the **Mount** command to mount initialized media.

▶ 1. Make sure the medium is initialized.
   See 8.3 Initializing a Tape or Disk for Recording (page 284) .

2. Insert the tape or disk into the drive and set it on line.
   See the manufacturer's documentation.

3. Choose **Commands > Mount**.

IRIS scans the medium to determine what kind of medium has been mounted, and checks the system disk to see whether there is a corresponding tape log.

The **Status** field reports Scanning during the Mount operation.

IRIS reports its findings in the **Filter Summary** field as one of:

- `40*/250 Files 11 M/ 202 M Bytes`
  For example, this indicates that the medium is a valid IRIS archive medium with a complete log. The log list shows the first few product files that are on the medium (if any). In this example, there are 250 files and 202 MB on the medium. The log list is currently displaying 40 files that occupy 11 MB (* means there are more that could be displayed).
- `Incomplete Log`
  This indicates a valid IRIS tape but an invalid or incomplete tape log on disk. The tape log shows the first few files on the tape to the best of its ability. You may want to perform an inventory if a Retrieve operation is to be done.

After a successful **Mount** operation, IRIS displays the archive ID, which should be checked against the label on the medium.

If a medium of unknown format or a blank tape or disk is detected during Mount, IRIS does not mount the medium and the **Status** field shows `NonIris`.

8.5 Recording Data

The **Record** command starts the recording process. Output requests made from the **Product Output** menu for the selected archive drive are written at the beginning of a blank tape or added to the end of an IRIS tape. The archive medium can be left mounted so that new files can be automatically written to the archive without having to enter the **Archive** menu. This continues until the medium is full.

Retrieve, **Unmount**, and **Inventory** commands can be issued during the record operation without first issuing the **Stop** command. A side effect of recording on an IRIS tape is that a new tape log is created when IRIS scans to find the last tape file. If you record on a tape that has an `Incomplete Tape Log`, a new tape log is generated. There is no need to perform an inventory in this case.

1. Enter the **Product Output** menu and select which product files to send to the archive drive. Be sure to specify the correct drive when you do this. You can make either individual send requests or use the header line commands to specify that all future products be sent to the archive.

2. Put a partially filled or a freshly initialized medium into the appropriate archive drive. Make sure that the medium is write-enabled.

3. Select **Menus > Archive Menu** and select the appropriate archive drive.
   Check that **Status** field indicates **Idle** or **Released**.

4. Select **Commands > Mount**.

5. Check the **Filter Summary** to make sure the medium mounted successfully, and check the **Archive ID** field and log list to make sure that this is the medium that you want to write.

6. Select **Commands > Record**.
   When recording is complete, the **Commands** field shows `Record Done` and the **Status** field shows `Idle`.
   The log contains entries for the new files that have been added.
   You can issues further commands (such as **Unmount**), but many users leave the medium is left in this state so that future requests for archive from the **Product Output** menu are fulfilled automatically.

7. If the medium is full before recording is complete, the tape rewinds to BOT and the system-wide message `TAPE FULL ON DRIVE (drive id)` is flashed.

   a. Enter the **Archive** menu and unmount the tape.

   b. Mount a new tape.

   c. To restart the recording operation where it left off, select **Commands > Record**.

   Products are queued to be sent to the drive and are sent as soon as you select the record command on the fresh media. Thus if the media before these queued files are deleted by the normal watchdog process, then there is no loss of data on the archive. The time depends on the size of your IRIS system disk but is typically several hours.

8. To stop recording manually, select **Commands > Stop**.
   IRIS waits for the next command, which could be to resume the record process. The log reflects the new files that were added before the `Stop` command was issued.

8.6 Creating and Printing a Log

> Disk file inventories are not created for LDA devices. The inventory for these is made on-the-fly.

The **Inventory** command is invoked when there is no tape log or when the tape log becomes damaged. If the latter case, the Archive ID field displays `Incomplete Log`.

For optical disks, logs are created automatically when a disk is mounted.

Printing is done on a text printer rather than a color printer, depending on how the system manager has configured your system. A tape does not need to be mounted to print the tape log.

8.6.1 Creating a Tape Log

1. Select **Commands > Inventory**.
   IRIS scans the tape and generate a new log. This can take some time.
   Note that when IRIS appends files to an existing IRIS tape, a new tape log is automatically generated.

2. When the tape is full, print a hardcopy tape log and store it in a notebook or file near where the tapes are stored. (You should also set the write protect.)
   This makes it easier to determine what product files are on a tape when data are retrieved.

8.6.2 Printing a Tape or Disk Log

▶ 1. From the operating system prompt, change directories to the location of the logs:

```
$ cd $IRIS_INV_TAPE
```

2. To see a list of the available tape logs in this directory, type:

```
$ ls
```

A separate file is listed for each archive log, identified by the archive ID corresponding to the label on the tape or disk.

3. Turn on the default printer for your system and type (for example):

```
$ lp SIG_88_03_08_S9C.LIS
```

The archive ID **SIG_88_03_08_S9C** is an example.

8.7 Retrieving Product Files from an Archive

You can use the **Archive** menu to retrieve product files from tape and put them back on disk where they are available for display or, in the case of the RAW product, for subsequent product generation.

The **Retrieve** command takes product files from archive and stores them on the system disk. Retrieve requires that you first tag all files to be retrieved. Files can be tagged even while **Retrieve** or other operations are in progress.

▶ 1. Because **Retrieve** relies heavily on the log, perform an inventory for IRIS tapes with incomplete logs.
   See 8.6 Creating and Printing a Log (page 287).

> ℹ️ **Retrieve** works if an inventory is not done, but only those files that are listed in the tape log can be retrieved.

2. Determine which tape or disk has the data that you want. The best way to do this is to maintain a book of printed logs. Because IRIS media can contain a wide variety of products, a simple label is not usually sufficient to describe what is on the medium.

3. Load the medium into the drive and set the drive on line.

> **i** Set the write protection on to avoid accidentally writing to the tape by yourself or another operator.

4. Enter the **Archive** menu and choose **Commands > Mount**.
   Check the information in the log to make sure it is consistent with what you thought was on the medium.

5. To tag files for retrieving, select the files that you want to tag by doing one of the following:
   • Left-click the mouse over an entry in the log to select it, or press **CTRL**+click to select more than one entry.
   • Click and drag the mouse to select a group of entries in the log, or press **CTRL**, then click and drag the mouse to select multiple groups of entries.
   • Use the **Filter** menu to produce a list of the files that you want to retrieve.

6. Select **Retrieve** from one of the following menus:
   • If you selected the entries with the mouse, position the mouse cursor over the tape log and right-click the mouse.
   • If you selected the entries from the **Filter** menu, select **Commands**.

   If you change your mind, untag the file by selecting **Cancel**.
   When a product is selected for **Retrieve**, an **r** appears next to its file number.
   After a file is successfully retrieved, or if it is on disk already, a **d** is displayed next to the corresponding file number.

> **i** Files that are already on disk cannot be tagged for **Retrieve**.

7. Select **Commands > Retrieve**.
   IRIS positions the tape and copies the tagged files to the disk. A **d** appears next to a file after it is copied to disk.
   You can continue to tag or untag files even while **Retrieve** is operating.
   You may also leave the **Archive** menu, and the **Retrieve** process continues its job.
   For example, you can go to the **Product Output** menu and request output from some of the products that you have restored to disk.

8.8 Stopping an Archive Operation

Use **Commands > Stop** to stop the current archive operation.

Issuing a **Stop** command during a **Record** or **Retrieve** operation does not damage any files.

The **Stop** command has different effects depending on the operation being stopped:

During Record
>
> **Stop** allows the file currently being written to be completed.
> If **Record** is re-issued, IRIS resumes recording at the end of the last file.

During Retrieve
>
> **Stop** finishes retrieving the last file, then stops retrieving.
> If **Retrieve** is re-issued, the retrieve process continues where it left off.

During Inventory
>
> **Stop** causes the inventory to halt. The resulting log is incomplete and the filter summary shows `Incomplete Log`.
> **Retrieve** operations that are limited to the part of the log that was completed can still be done. If the **Inventory** command is reissued, the inventory is started from the beginning of the tape.

During Mount
>
> **Stop** cancels the **Mount** operation and rewinds the tape.

During Unmount
>
> **Stop** has no effect.

8.9 Unmounting a Tape or Disk

The **Unmount** command causes IRIS to rewind a tape and take the tape or optical disk drive off line. The log information is cleared from the screen. This command must be invoked before a new tape can be mounted.

1. Select **Commands > Unmount**.

2. Wait for the tape to completely rewind, then remove it from the tape drive.
   If the media is a DVD, the contents of the temporary buffer used for writing the DVD are flushed to the DVD, that is, you do not lose the data in the buffer.
   This may delay the unmounting step by 20 seconds. If you quit IRIS (`qiris`) without unmounting the DVD, the buffer is flushed in the same way (delaying the `qiris`).

9. Managing Ingest Files

9.1 Managing Ingest Files

Ingest files are the data collected from the signal processor each time a task runs. An ingest file is not the same as a RAW data product, which stores ingest data in a compressed format.

Ingest files are cataloged by their associated tasks and stored on disk in an ingest summary file. IRIS uses an automatic Watchdog process to delete old ingest files.

The **Ingest Summary** menu provides a list of the ingest files currently on disk. Use the top part of the menu to choose the files that you want to display. The bottom part of the menu lists the ingest files.
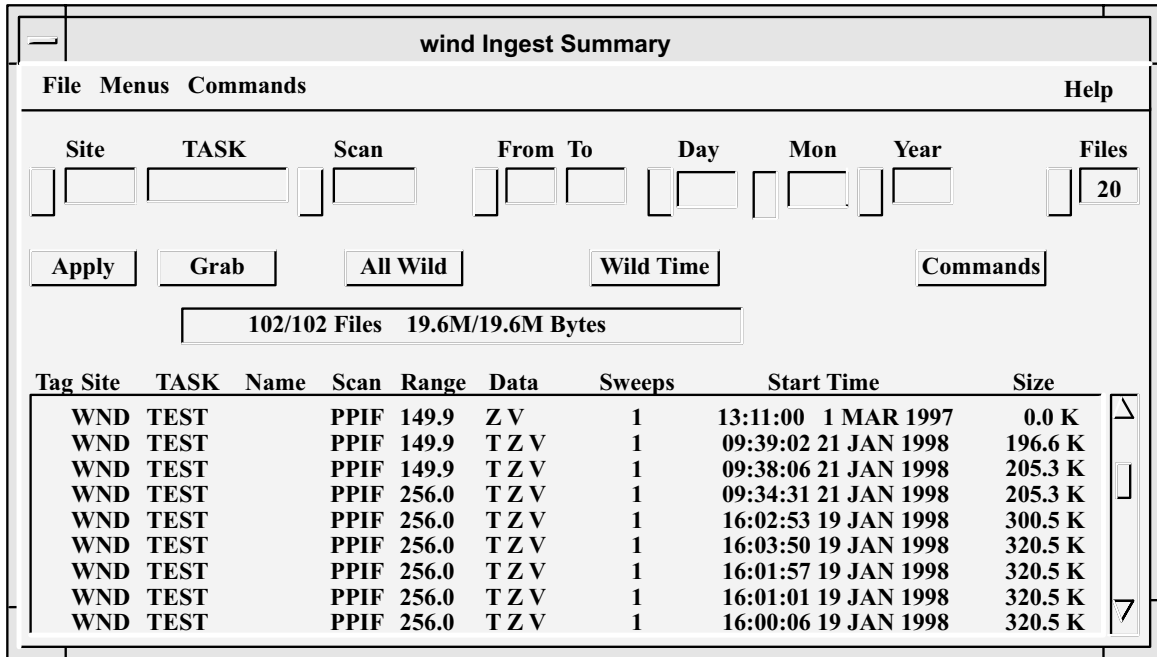
Tag Site	TASK	Name	Scan	Range	Data	Sweeps	Start Time	Size
WND TEST			PPIF	149.9	Z V	1	13:11:00 1 MAR 1997	0.0 K
WND TEST			PPIF	149.9	T Z V	1	09:39:02 21 JAN 1998	196.6 K
WND TEST			PPIF	149.9	T Z V	1	09:38:06 21 JAN 1998	205.3 K
WND TEST			PPIF	256.0	T Z V	1	09:34:31 21 JAN 1998	205.3 K
WND TEST			PPIF	256.0	T Z V	1	16:02:53 19 JAN 1998	300.5 K
WND TEST			PPIF	256.0	T Z V	1	16:03:50 19 JAN 1998	320.5 K
WND TEST			PPIF	256.0	T Z V	1	16:01:57 19 JAN 1998	320.5 K
WND TEST			PPIF	256.0	T Z V	1	16:01:01 19 JAN 1998	320.5 K
WND TEST			PPIF	256.0	T Z V	1	16:00:06 19 JAN 1998	320.5 K

Figure 98   **Ingest Summary** Menu

1. Select **Menus > Ingest Summary**.
   The **Ingest Summary** list in the bottom pane is sorted by time, with the newest files listed first.

   Tag
   Displays k for keep, d for delete, or blank if no flag is set.

   Site
   Shows the site ID for the radar site from which data were gathered.

   TASK
   Contains the name of the task.

   Scan
   Contains type of scan used to collect the data, such as **RHI** or **PPI**.

   Range
   Contains the value of the **Max Range** field, taken from the task configuration.

   Data
   Lists the data parameters (Z, T, V, W) for the task.

   Sweeps
   Shows the number of sweeps that were collected by the task. If the scan was halted prematurely, this column displays the partially completed task as N/M, even if the task was halted during the final sweep.

   Start Time
   Shows the time that the task started to collect data.

   Size
   Shows the total amount of disk space occupied by the ingest files for this run of the task.

2. Select an entry in the list and position the mouse cursor over any of the fields to show the following options:
   - **Keep** tags all selected files with the keep flag so they are not deleted by the Watchdog.
     This is necessary if you want to save ingest files, including those that have been restored from tape, for non-real time product generation in an analysis mode. For example, you might want to make a product tomorrow from today's ingest data. See 9.2 Tagging and Untagging Ingest Files (page 295).
   - **Clutter** tags a selected file as a clutter map.
   - **Un-Tag** removes delete or keep flags from the selected files.
   - **Delete** tags all selected files with the delete flag to free more disk space for new ingest files.

3. Use the **Filter** menu in the top pane to narrow the list down to those files that are of interest by site, task, scan, or time.

   Site

   Enter a site ID to select data from only one radar site or enter the wildcard character to select data from all sites.
   Select **Site** to show a list of known sites. The contents of this list depend on the sites that are configured in the **Setup** utility.

   TASK

   Enter the name of a task to filter for products generated by a specific task.
   You can include wildcard characters in the task name.
   For example, entering a product name of `*PPI*_?` selects all hybrid tasks that have `PPI` somewhere in their names.

   Scan

   Enter a scan type in this field to select products of only one type, enter the wildcard character to select products of all types, or select from a list of types.

   From, **To**

   Enter a range of hours to filter products by time of day.
   You can also select from a list of ranges.

   Day, **Month**, **Year**

   Enter a day, month. or year to filter products by date.
   You can also select from a list of values.

   Files

   Control how many files to include in the list.

   Apply

   Select **Apply** to update the product list, based on your selection criteria.

   Grab

   If you have selected a product, select **Grab** to insert the product's site, type, product name, task, and date information in the filter fields.

   All Wild

   Select **All Wild** to return all the fields to the wildcard character.

   Wild Time

   Select **Wild Time** to change only the hours, month, day, and year fields to the wildcard character.

   Commands

   Apply tags to the ingest files selected by the **Filter** menu:
   - **Keep** tags all selected files with the keep flag so they are not deleted by the Watchdog.
     This is necessary if you want to save ingest files, including those that have been restored from tape, for non-real time product generation in an analysis mode. For example, you might want to make a product tomorrow from today's ingest data.
     See 9.2 Tagging and Untagging Ingest Files (page 295).
   - **Clutter** tags a selected file as a clutter map.
   - **Un-Tag** removes delete or keep flags from the selected files.
   - **Delete** tags all selected files with the delete flag to free more disk space for new ingest files.

9.2 Tagging and Untagging Ingest Files

> ⚠️ **CAUTION!** Deleting files may interfere with subsequent product generation. Be sure that a file is not required for product generation before you delete it. The only circumstance that may require manual deletion is when you need to make room when restoring a RAW file from tape.

1. Select one or more files from the list. There are several ways to make your selection:
   - Left-click the mouse over an entry in the list to select it, or press **CTRL**+click to select more than one entry in the list.
   - Click and drag the mouse to select a group of entries in the list, or press **CTRL**, then click and drag the mouse to select multiple groups of entries.
   - Use the **Filter** menu to produce a list of the ingest files that you want to tag.

2. Choose **Keep**, **Delete**, or **Un-Tag** from one of the following menus:
   - If you selected the entries with the mouse, position the mouse cursor over the ingest file list and right-click the mouse.
   - If you selected the entries from the **Filter** menu, select **Commands**.

   The files are marked with a **k** for keeping, a **d** for deleting, or the tags are removed, depending on the command you chose.

10. Choosing Overlay Files

10.1 Overlay Overview

You can use the IRIS **Overlay** utility to draw overlays or maps for display on top of other IRIS/Open products. Overlays are used for product output and the real-time display.

An overlay can consist of the following:

- Geographical and political boundaries displayed with or without latitude and longitude lines or range rings.
- Text strings to label areas of interest, such as cities.
- Bitmap Icons that can be constructed by the user and applied to the overlay. Icons can represent any feature, for example, airports, train stations, and similar.
- The lines, text and icons mentioned above can be separated into different layers within the overlay file. Then at display time, either all or only a subset of these layers may be displayed giving yielding overlays that appear different based on which layers are active. Layers can be drawn in different colors.
- Underlays are filled regions of color displayed under the radar images where there is no weather data. Underlays are typically used to indicate areas of water.

Each overlay is defined in an ASCII file using a connect-the-dots approach (sometimes called a vector approach). This allows overlays to be drawn to any scale factor.

You use the **Overlay** utility to create and modify your own overlays in a window on your workstation.

Typically, you get the basic map (coastlines, borders, rivers) from Vaisala and then customize it to meet your needs. It is often convenient to make separate layers for different interests (river catchments, airports). IRIS supports up to 20 overlay files.

IRIS overlays are ASCII files so you can edit them with any text editor (`vi`, or `emacs` which is distributed with IRIS).

10.2 Invoking Overlay

▶ 1. Invoke the **Overlay** utility by typing: `overlay&`
Overlay starts as a default with an empty window on your workstation.
Settings from your previous session are still valid when you restart.

2. Open an overlay (*.ovr*) file from the */usr/sigmet/config/overlay* directory.
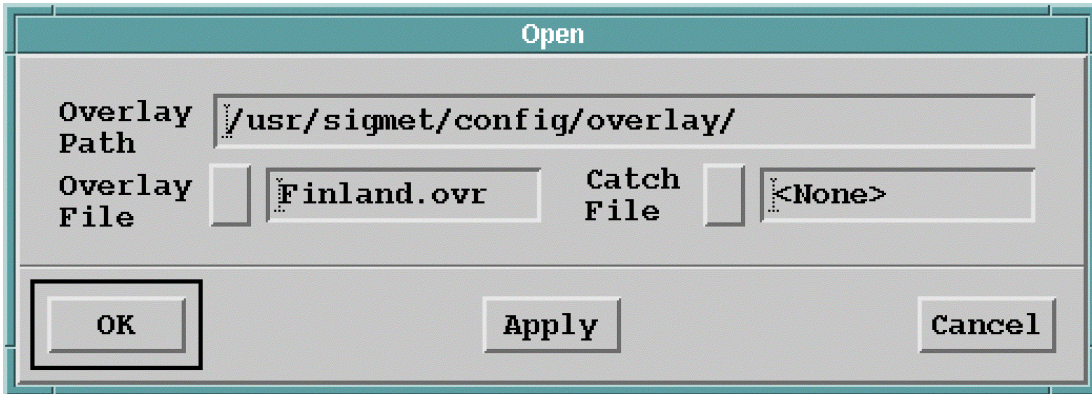
3. To see the latitude and longitude of your cursor, select **File > Display Lat Lon**.
   You can move your cursor around your map with the mouse. To find a precise point, use
   the arrow buttons of the keyboard to move pixel by pixel.

4. To define the mode to show latitude and longitude, select **File > Lat Lon Units**.
   You can display the values as degrees and decimals or degrees, minutes and their
   decimals.

5. To define the map display, select **File > Options**.
    The **Options** window is shown.

a. In the top pane, define the following:

   Display size in km
      Defines the horizontal size of the map in km.

   North/East offset in km
      Move around your original map and create subareas.

   Layers to draw
      Defines if you want to work with all of your map layers or only a subset.

b. In the bottom pane, define the following:

   Fill points
      Indicate the areas, typically lakes, you want to be colored with a certain color.

   Range Rings
      Centered at a predefined point, typically a radar site

   Protected
      The areas you use with the WARN product

   Lat Lon
      Displays a latitude-longitude grid

NoUnderlay

Do not use an underlay.

c. Select the map projection.

Typically, **Azimuthal** is used in one–radar applications, **Mercator** in composites in the tropical areas and **Polar Stereographic** in the composites near the poles.

6. To create or modify a layer, select **File > Add/Modify Layer**.
    Everything you add to an overlay goes to a layer. You can add or delete text, icons and fill points. You can also select a "pointer" which gives you position of you cursor.
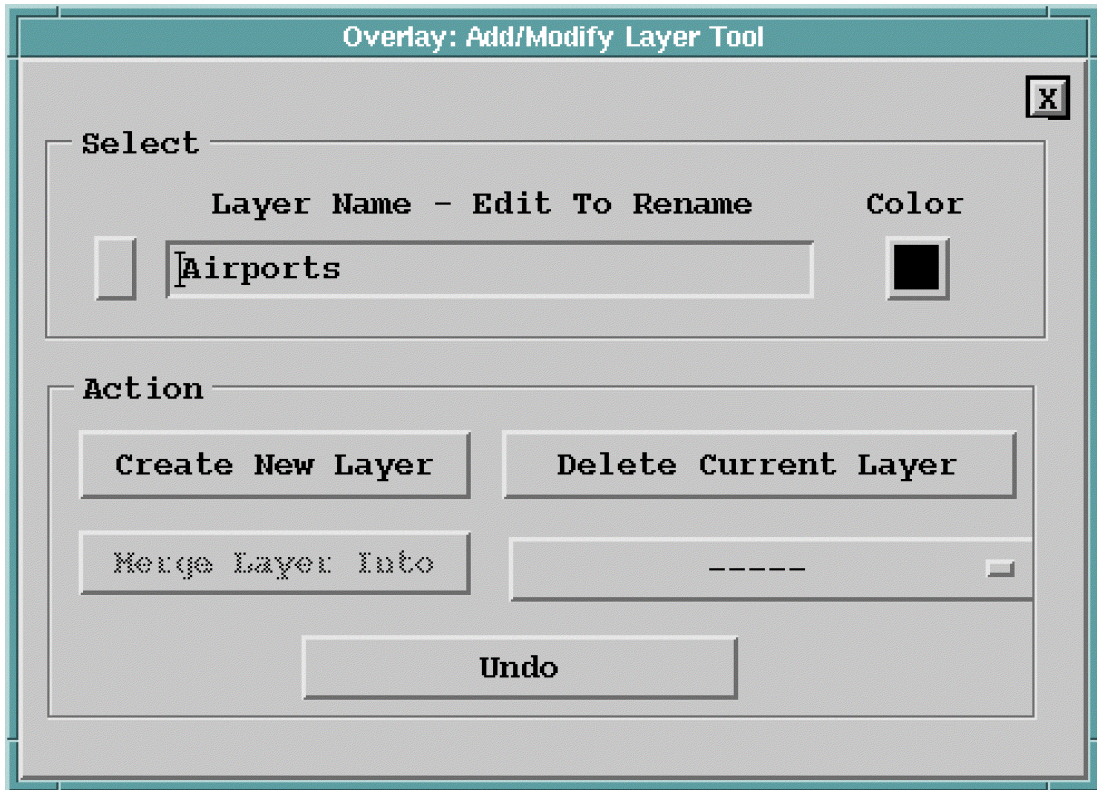
a. Assign each layer a color from a predefined set, shown in a box next to the layer name.
    To define more colors, use the **Color** utility.

b. To add content to a new or existing layer, select **Create New Layer**.
    The **Add/Delete Tool** appears.
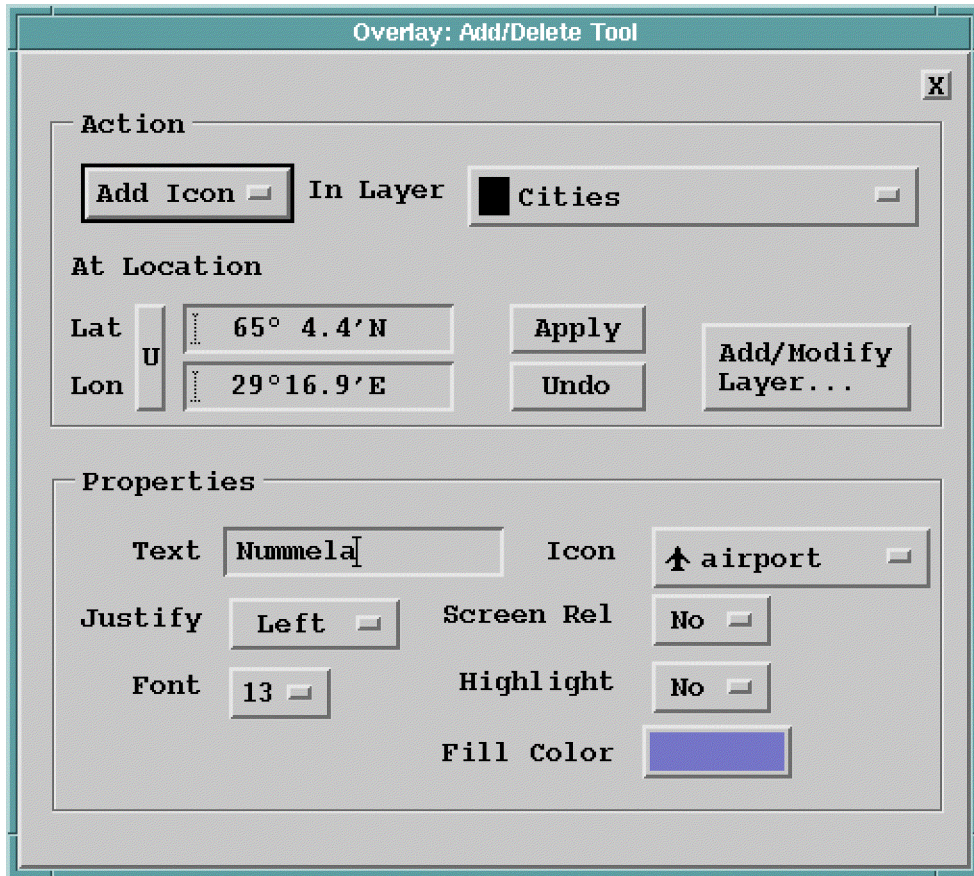
- To add icons, you must include a text next to the icon.

When you use icons, Vaisala recommends that you justify them right or left, because the icon goes to the point and the text should be next to the icon. You can highlight the text to give it a colored background.
- When adding text, most users want to justify text to center, so that the point shown with latitude and longitude in the upper column is under the middle character of your text.

c. To merge 2 layers, in the **Add/Modify Layer** window, select target layer and **– – –**.

d. To prevent the filled color from leaking around the edge of the overlay limits, use the **RANGE** command to tell **Overlay** how far you have drawn your overlay.

e. Define your underlay fill points.
Fill points are used to make an underlay. An underlay can contain up to 400 fill points. **Overlay** starts from each of the specified fill points and fills all contiguous areas, stopping at overlay lines.
You can enter many fill points for the same filled region. In fact, this is required for a coastline with deep bays, so that the region fills correctly at all scales. Many overlays contain coastline and use a blue underlay for ocean.

f. To delete items from a map, select a **Kill** option.
When you select a **Kill** option, your cursor becomes a killer tool and you can delete items from the map by pointing and clicking them. There is a separate kill tool for text, icons and fill points to minimize the risk you kill something you didn't want to kill. If you can't kill what you want, you are probably working with a wrong layer.

7. Select **Header** to check the overlay file header.

See .
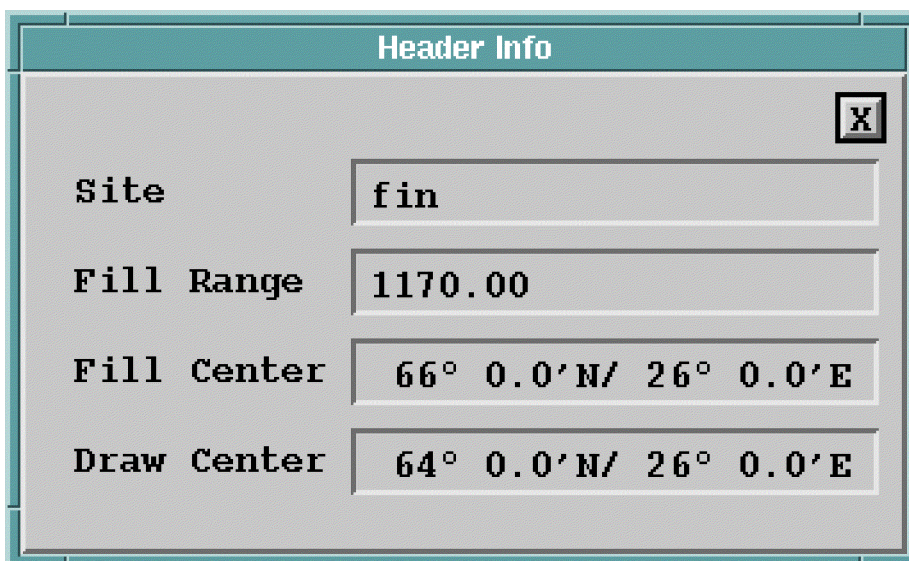
10.3 Overlay Header

Absolute Coordinates

If the overlay is defined in absolute coordinates, all points in the overlays are given in terms of a longitude, latitude coordinate pair. This sets an absolute position for each point.

The following header commands must be included at the top of the overlay file:

Table 22   Overlay Latitude and Longitude Header Commands

Command	Description
PROJECTION latlon	Informs system overlay is in absolute coordinates.
DRAW_CENTER lon lat	Tells **Overlay** where to center the image. Generally this is the either the coordinates of the radar site, or a nearby special area of interest such as an airport.
FILL_CENTER lon lat	Specifies center of fill area. Typically this is the center of the array of data points you have generated.
FILL_RANGE range	Sets the maximum range that the underlay covers, in km. **Overlay** draws a circle at this distance from the fill center before filling the underlay.

Relative Coordinates

If the overlay is defined in relative coordinates, all points are defined in terms of map units. Map units are any arbitrary unit of measure that is E/W and N/S of an origin (0,0), which can be anywhere on the map.

Often millimeters from graph paper are used to designate the points on a map. The map must be in azimuthal equidistant projection, with the specified reference point. A scale factor specifies the conversion between these map units and distance on the surface of the Earth.

The following header commands must be included at the top of the overlay file:

Table 23   Overlay Map Unit Header Commands

Command	Description
PROJECTION aed	Informs system overlay is in relative coordinates.
SCALE scale	Sets the scale factor in km/map unit. The **SCALE n** command must be used to specify how many kilometers each map unit specifies. For example, if a **RADAR 100 200** and a **SCALE 0.5** command are used, then if later a point is specified as **POINT 110 200**, this point is exactly 5 km due east of the reference position.

Command	Description
`RADAR x y`	Sets the projection point in map units and specifies the location of the reference point in the map units.
`REFERENCE lon lat`	Sets the projection point in absolute coordinates.
`DRAW_CENTER lon lat`	Same as `latlon` projection.
`FILL_CENTER lon lat`	Same as `latlon` projection.
`FILL_RANGE range`	Same as `latlon` projection.

For backwards compatibility, the old **MIDDLE** command still has the meaning of setting both the draw center and the fill center.

10.4 Assigning an Overlay to a Radar Site

The IRIS system manager associates a default overlay file for each site connected to the system. You can assign a different overlay to the site.

1. In the **Overlay** menu, highlight the site you want to change.
   A list of available overlays is shown.
   The following figure shows an example with 2 overlays for the **HEL** site: **HELSINKI** and **KARTTA**.
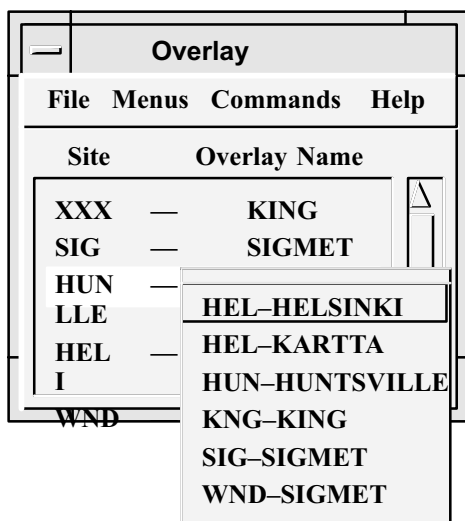
Figure 99   Site Overlay Example

2. Select an overlay from the list.
   The menu also shows the site IDs for which the overlay is valid.
   If you try to display a product with an invalid overlay, IRIS displays the message,
   `Overlay origin more than 1000 km from radar location` (1000 km = 621.4 mi). If you get this message, check that the correct overlay is assigned to the radar site.

Appendix A. Basic Radar Meteorology

A.1 Introduction to Radar Meteorology

A weather radar sends pulses of microwaves and measures, how much is reflected by targets in the atmosphere. Targets are mainly raindrops, hail, or snow.

A radar antenna focuses the microwaves to a beam of 1°. The beam grows wider in distance: at distance of 200 km (124.3 mi) it is already 4 km (2.5 mi) wide.

Microwaves travel at the speed of light. Measuring the time when the reflection comes back to antenna, we can count how far the target is.

When a microwave pulse meets some raindrops, very small amount is reflected (scattered) and most of the waves go on. This is because typically raindrops are size of 1 mm (0.04 in) distance of 10 cm (3.94 in) from each other. This means we can see "rain behind the rain". This also means a radar must be very sensitive: we send about 250 kW and receive 10^{-13} W.

We can move the antenna up and down (elevation) and round in horizontal plane (azimuth). Typically, in a measurement task, we take a low elevation and measure in all azimuths, then increase the elevation a little, and measure in all azimuths again, and so on, so we get data from all directions at different ranges. This pack of data is called a polar volume, "polar" because it is measured in polar coordinates with the antenna in the center.

The radar alternates between sending and receiving at "pulse repetition frequency", (PRF). That is typically between 100 and 1000 times a second. When the antenna is moving in azimuth direction all the time, each pulse is sent in slightly different direction. In signal processing, we process the data from same distance from different pulses. These are called samples. Slices in distance which are processed are called bins. In the image in next page, value of each bin is processed from four samples. The distance between samples is determined by antenna speed and PRF. To get reasonable values, usually 32 … 128 samples are processed together.

Figure 100 Range Bin Geometry

A.2 Reflectivity

It is convenient to say things in atmosphere reflect the radiation, but backscattering is the accurate physical term. It is measured in **Z** (radar reflectivity factor) and often expressed in **dBZ** (radar reflectivity).

The famous radar equation from classical optics says

$$Z = \frac{\mathrm{Pr}^2}{LCK}$$

- **P** is the measured average power (in Watts) of several samples at the radar
- **r** is the range to the bin
- **L** is attenuation
- **C** is radar (hardware) constant
- **K** is the refractive index and depends on the dielectric properties of the particle

For meteorologists

$$\mathrm{Re}\,flectivity = \frac{(Watts\_\mathrm{Re}\,ceived \times Dis\tan ce\_Squared)}{(Attenuation \times Hardware\_Cons\tan t \times Rain\_or\_Snow\_Cons\tan t)}$$

For cloud physics, **Z** is a sum $\Sigma N_i D_i^6$ where N_i is the number of particles with diameter D_i per unit volume in the atmosphere. That means, that one droplet with diameter 4 mm (0.16 in) gives 4096 times as much energy as a 1 mm (0.04 in) droplet. And that we can't know if there is one droplet of 2 mm (0.08 in) or 64 droplets of 1 mm (0.04 in).

Z varies between 0.001 and 10,000,000. To get understandable numbers, we use a decibel scale:

$$dBZ = 10\log\frac{\left(Z mm^6 mm^{-3}\right)}{\left(1 mm^6 mm^{-3}\right)}$$

The following figure shows typical values for various phenomena in the atmosphere. You can see that reflectivity strength alone is not enough for target identification.

Figure 101   dBZ Values for Various Phenomena

For hydrology, we need an equation to combine radar reflectivity to rainfall rate. All of these equations are empirical and approximate, since **Z** is proportional to **D6** and precipitation rate **R** (mm/h) is proportional to **D3.7**. Also, we must assume something about drop sizes. For 64 drops of diameter 1 mm (0.04 in) or one drop, diameter 2 mm (0.08 in), **Z** is same, **R** is not. A good first guess is the classical Marshall Palmer equation:

$Z = 200R^{1.6}$ which equals to $R = (Z/200)^{0.625}$

A.3 Geometry

The Earth is spherical. Microwaves do not follow the surface. Because of the atmospheric optics, they usually bend down a bit, at a radius of 4/3 of the Earth.

The following figure shows the height of some radar beams as function of distance from radar at different elevations.

Though nominal beam width is 1°, some energy is spread to other angles, too. These are sidelobes.

Figure 102   Beam Height vs. Range

A.4 Problems in Interpreting Radar Images

Echoes not related to rain:

- Reflection of microwaves from ground or water surface, radar mirage
- Birds, insects, chaff from airplanes
- Solid targets such as masts and buildings
- Other microwave sources such as the sun and other radars
- Second trip echoes

Echo intensity differences not related to differences in precipitation intensity:

- Calibration issues: nowadays electrical calibration can effect 2 to 3 dB.
- Beam blocking: Buildings near the antenna can shade partially or totally echoes from certain direction
- Attenuation in rain: echoes behind heavy rain can disappear especially if wavelength is small. Attenuation correction can compensate if echo is too weak but cannot help if there is no echo.
- Water phase: ice, melting ice and water scatter microwaves differently. Bright Band can be 7 dBZ
- Drop size: especially melting hail. 2 mm (0.08 in) droplet = 64 of 1 mm (0.04 in) droplets

Even if the measurement is OK, there is not a 100 % correlation with a gage:

- Far from antenna the beam is rather high. It can overshoot a precipitating cloud. Or the rain can evaporate below the measurement height. Orographical enhancement can occur.
- In cold winters, the difference can be 10 dB between measurement height and surface. Vertical reflectivity profile correction can help to fix quantitative errors, but if there is no measurement the correction does not help.

A.5 Doppler Wind Measurements

In comparison to non-Doppler radar, Doppler radar can provide clutter cancellation and wind products.

A Doppler radar measures the speed of the reflector by studying how the phase of microwaves compared to the sent pulse changes between samples.

- A non-moving target has no phase shift.
- A target whose movement has no component in direction of the radar beam, has no phase shift.
- A phase shift of exactly 2p looks like no phase shift.
- A phase shift of 2p+d looks like phase shift d.
- A phase shift of n*2p+d looks like phase shift d. This is called aliasing.

To get good Doppler measurements, many samples must be processed together (64 is good). This sets limits to antenna speed.

The maximum speed which can be measured unambiguously is

$$V_{\max} = \frac{PRF \times \lambda}{4}$$

where λ (lambda) denotes wave length.

Typical pulse repetition frequencies (PRFs) in weather radars are 250 ... 1200 Hz.

Because the PRF also determines the maximum unambiguous range, it is not possible to just use 1200 Hz and get nice winds.

$$R_{\max} = \frac{c}{2 \times PRF} = \frac{Speed\ of\ Light}{2 \times\ PRF}$$

This leads to the Doppler dilemma: we can measure either high speeds or far away from the radar - and we want both. Dual PRF techniques with single PRF algorithms and hybrid tasks employing different PRF for different elevations.

A.6 Clutter Cancellation

Clutter is defined as echoes from hills, buildings, masts, sea, and noise as marks in the image caused by electronics of the equipment. The goal of clutter cancellation is to remove clutter without destroying rain data.

To illustrate a Doppler filter we study speed spectra of a bin. That is, Doppler speed of each sample, horizontal axis being speed (towards or away from radar, 0 in the middle). If all samples indicate no speed we know the target is not moving. However, even antenna movements gives some speed. And there are cases when it is raining at the hills.

The next image shows a Doppler spectra from a bin, which contains clutter (speed near 0) and rain moving toward the radar (big hump on the left) as well as some noise (small humps everywhere, seen alone on the right). To find the right settings for the Doppler filter, we can move the red lines closer to each other (less rain data is destroyed) or further from each other (more clutter is canceled).

Figure 103   Typical Spectrum Plot

Appendix B. Radial Velocity Correction

B.1 Radial Velocity Correction

The motion of a ship or airborne platform induces a radial velocity error. IRIS corrects for this error by accounting for the platform motion.

Ship motion measurements are collected from the antenna controller and the radial velocity correction is optionally applied to the ingest files so that all subsequent products and the real time display are corrected. Note that the radial velocity correction and all of the platform motion parameters are archived as part of the extended ray header.

B.1.1 Motion Parameters and Coordinate Transformations

For systems with the Vaisala RCP8, see the theoretical derivation of the stabilization in *Radar Control Processor RCP8 User Guide*.

Table 24   Basic Variables of Ship Motion Sensing

Symbol	Definition
H	Heading measured CW from true north
θ	H - π/2
`#thetaoverdot#`	Rate of change of heading
P	Pitch measured relative to horizon + for bow down
`#Poverdot#`	Rate of change of pitch
R	Roll measured relative to the horizon + for port side down
`#Roverdot#`	Rate of change of roll
u	Eastward velocity
v	Northward velocity
w	Upward velocity

Given a vector **x** in earth coordinates and the corresponding vector **x**`#prime#` in ship coordinates,

$$x' = Ax \quad \text{and } x = A^{-1}x'$$

where the transformation matrices are:

$$A = \begin{bmatrix} \cos R \cos H - \sin R \sin P \sin H & - \cos R \sin H - \sin R \sin P \cos H & - \sin R \cos P \\ \cos P \sin H & \cos P \cos H & - \sin P \\ \sin R \cos H + \cos R \sin P \sin H & - \sin R \sin H + \cos R \sin P \cos H & \cos R \cos P \end{bmatrix}$$

$$A^{-1} = \begin{bmatrix} \cos R \cos H - \sin R \sin P \sin H & \cos P \sin H & \sin R \cos H + \cos R \sin P \sin H \\ - \cos R \sin H - \sin R \sin P \cos H & \cos P \cos H & - \sin R \sin H + \cos R \sin P \cos H \\ - \sin R \cos P & - \sin P & \cos R \cos P \end{bmatrix}$$

Note that A^{-1} is the transpose of A. Where the vectors are defined as:

$$x = \begin{bmatrix} East \\ North \\ Up \end{bmatrix} \qquad x' = \begin{bmatrix} Towards\ right \\ Towards\ bow \\ Towards\ zenith \end{bmatrix}$$

B.1.2 Radial Velocity Correction

The scalar radial velocity of the scatterers V_R can be expressed as:

$$V_R = V_{R\,meas} - V_{R\,ship}$$

where $V_{R\,meas}$ is the Doppler velocity measured by the signal processor and $V_{R\,ship}$ is the radial velocity correction for ship motion. The radial velocity correction is found by taking the dot product of the antenna vector x in earth coordinates with the velocity of the antenna in earth coordinates y_A, as follows:

$$V_{R\,ship} = x \cdot \dot{y}_A$$

The antenna vector is a unit vector pointing in the direction of the antenna which depends on the true (earth relative) azimuth and elevation of the antenna, as follows:

$$x = \begin{bmatrix} \sin AZ \ \cos EL \\ \cos AZ \ \cos EL \\ \sin EL \end{bmatrix} \qquad AZ and\ EL\ are\ earth\ relative$$

These are known from the antenna controller. Thus, the problem is to determine the y_A, the velocity of the antenna.

The velocity correction used in IRIS allows for the following types of ship motion sensing:

Inertial Navigation Unit (INU)
   (*INU Case*) The INU is typically placed close to the antenna. The INU reports pitch, roll, and heading (and their rates of change), as well as the east, north and vertical motion (u, v, w). In this case, the INU is used as the ship reference point.

Gyro System with GPS

(*Gyro Case*) The gyros report pitch, roll, and heading (and their rates of change). Mean translational motion of the center of the ship (**u**, **v**, **w**) is obtained from a GPS or other navigational system. In this case the ship center of rotation is used as the ship reference point.

The interface either device is the RCP antenna control unit. The orientation and navigation information is passed to IRIS through the serial line connection to the antenna controller.
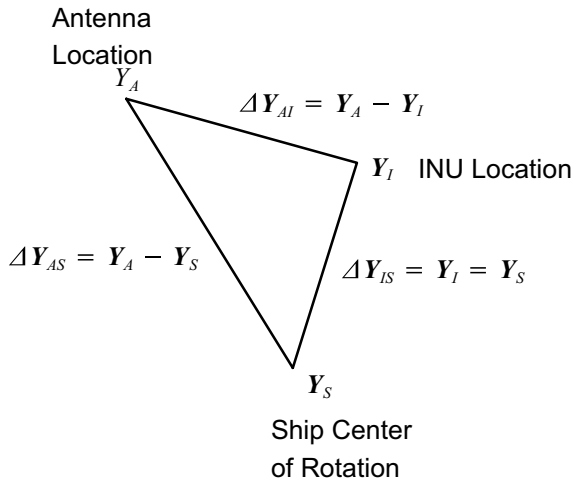
Figure 104   Vectors used for Velocity Correction In Earth Coordinates

In general, the INU is not located at either the ship center of rotation or the antenna. The location vectors are shown in the previous figure in earth coordinates. The position vector of the antenna in earth coordinates can be expressed as:

INU Case: $y_A = y_1 + (y_A - y_1)$

Gyro Case: $y_A = y_S + (y_A - y_S)$

so that,

INU Case: $y_A = y_1 + \Delta y_{AI}$

Gyro Case: $y_A = y_1 + \Delta y_{AS}$

where, $\Delta y_{AI} = y_A + y_I$ and $\Delta y_{AS} = y_A + y_S$

Expressing the Δ vectors in terms of the ship relative coordinates yields:

INU Case: $y_A = y_1 + A^{-1}\Delta y'_{AI}$

Gyro Case: $y_A = y_S + A^{-1}\Delta y'_{AS}$

Differentiating these two expressions yields the final expressions for the antenna velocity:

INU Case: $y_A = \dot{y}_I + \dot{A}^{-1}\Delta y'_{AI}$

Gyro Case: $y_A = \dot{y}_S + \dot{A}^{-1}\Delta y'_{AS}$

Note that since the INU, antenna and ship center coordinates are not changing relative to the ship coordinates,

$$\dot{y}_{AS} = \dot{y}_{AI} = 0$$

The components of the matrix A^{-1} are as follows:

$$\dot{A}^{-1} = \begin{bmatrix} \dot{A}_{11}^{-1} & \dot{A}_{12}^{-1} & \dot{A}_{13}^{-1} \\ \dot{A}_{21}^{-1} & \dot{A}_{22}^{-1} & \dot{A}_{23}^{-1} \\ \dot{A}_{31}^{-1} & \dot{A}_{32}^{-1} & \dot{A}_{33}^{-1} \end{bmatrix}$$

where:

$$\dot{A}_{11}^{-1} = -\dot{P}\sin P\cos\theta - \dot{\theta}\cos P\sin\theta$$

$$\dot{A}_{21}^{-1} = -\dot{P}\sin P\cos\theta - \dot{\theta}\cos P\cos\theta$$

$$\dot{A}_{31}^{-1} = -\dot{P}\cos$$

$$\dot{A}_{12}^{-1} = -\dot{R}\sin R\sin\theta + \dot{\theta}\cos R\cos\theta - \dot{R}\cos R\sin P\cos\theta - \dot{P}\sin R\cos P\cos\theta + \dot{\theta}\sin R\sin P\sin\theta$$

$$\dot{A}_{22}^{-1} = -\dot{R}\sin R\cos\theta - \dot{\theta}\cos R\sin\theta + \dot{R}\cos R\sin P\sin\theta + \dot{P}\sin R\cos P\sin\theta + \dot{\theta}\sin R\sin P\cos\theta$$

$$\dot{A}_{32}^{-1} = -\dot{R}\cos R\cos P + \dot{P}\sin R\sin P$$

$$\dot{A}_{13}^{-1} = +\dot{R}\cos R\sin\theta + \dot{\theta}\sin R\cos\theta - \dot{R}\sin R\sin P\cos\theta + \dot{P}\cos R\cos P\cos\theta - \dot{\theta}\cos R\sin P\sin\theta$$

$$\dot{A}_{23}^{-1} = +\dot{R}\cos R\cos\theta - \dot{\theta}\sin R\sin\theta + \dot{R}\sin R\sin P\sin\theta - \dot{P}\cos R\cos P\sin\theta - \dot{\theta}\cos R\sin P\cos\theta$$

$$\dot{A}_{33}^{-1} = -\dot{R}\sin R\cos P - \dot{P}\cos R\sin P$$

These equations show that the motion of the antenna can be expressed as a translational motion of a reference point (either at the INU or at the ship center of rotation), plus a rotational motion about a moment arm from the reference point to the antenna.

The translational motion of the reference point is:

As sensed by the INU for the *INU Case*: $\dot{y}_I = \begin{bmatrix} u \\ v \\ w \end{bmatrix}$

As sensed by the GPS for the *Gyro Case*: $\dot{y}_S = \begin{bmatrix} u \\ v \\ w \end{bmatrix}$

As a special case, if the INU and antenna are located at precisely the same position, then,

INU Case: $\Delta y'_{AI} = 0$ so that $\dot{y}_A = \dot{y}_I$

In this case, the antenna velocity is equal to the INU velocity since the moment arm is 0 length.

B.1.3 Aircraft Tail Radars

For aircraft tail radars, the math for the velocity correction is different.

Table 25   Velocity Correction for Aircraft Tail Radars

Parameter	Description
Azimuth	Because the axis of primary rotation is horizontal, azimuth is the angle relative to a perpendicular scan. Thus zero means scanning in a plane, positive means scanning a cone faced forwards, and negative means scanning in a cone faced backwards.
Elevation	Because the axis of primary rotation is horizontal, elevation is the angle of rotation about the axis of aircraft motion, with zero pointing out towards the left wing, and positive going up from there.
Recorded angles	Recorded angles are aircraft relative, and all angle syncing is done aircraft relative. It does not make sense to correct the angles to earth relative before recording.
Roll and pitch	Reversed sign on airplanes against ships. The computation of the A matrices defined for ships have the reversed sign. Note that since no other interpretation of the roll and pitch is done, this can otherwise be ignored.

The antenna vector relative to the aircraft is as follows:

$$x' = \begin{bmatrix} -\cos EL \ \cos AZ \\ \sin AZ \\ \sin EL \ \cos AZ \end{bmatrix} \quad AZ \ and \ EL \ are \ aircraft \ relative$$

Before computing the velocity correction dot product, we first must convert to earth relative, using the equation above:

$$x = A^{-1}x$$

After that, everything is the same as for ships.

B.1.4 Configuring Radial Velocity Correction

You must input the moment arm vector components in the **Setup** utility:

INU Case Δy'$_{AI!}$:

Δy'$_{AI1}$
　　Antenna distance forward of INU in deck plane.

Δy'$_{AI2}$
　　Antenna distance to port of INU in deck plane.

Δy'$_{AI3}$
　　Antenna distance upward from INU normal to deck plane.

Gyro Case Δy'$_{AS}$:

Δy'$_{AS1}$
　　Antenna distance forward of ship center in deck plane.

$\Delta y'_{AS2}$

Antenna distance to port of center line in deck plane.

$\Delta y'_{AS3}$

Antenna distance upward from ship center normal to deck plane.

Make the measurements to the antenna feed horn. The signs of the components must be strictly observed. The units are meters. Note that the questions in **Setup** refer to the INU case. If you have the *Gyro Case*, enter the corresponding moment arms relative to the ship center of rotation.

Setup has an additional question regarding the height offset of the INU. This is a fixed number which is added to the height reported by the INU. Since most INU's accurately report the height, this number should be set to zero. The value of this number does not affect the velocity correction.

- Specify the antenna format in **Setup**, RCP section.
- Specify the Extended Ray Header in **Setup**, ingest section.
- If you have a Vaisala RCP8, configure as described in *Radar Control Processor RCP8 User Guide*.
- Use the **TASK Configuration** menu to turn on and off the velocity correction for each task.

B.1.5 Testing Radial Velocity Correction

Vaisala provides several programs for testing the velocity correction feature.

B.1.5.1 Using Antenna Utility to Test Velocity Correction

The **Antenna** utility displays the navigation information that is reported back from the antenna controller as well as the velocity of the antenna Δy_A and the velocity correction.

Use the **Antenna** utility should to verify the following:

- The signs of all angles (pitch, roll, heading).
- The signs of angular velocities (pitch, roll, heading).
- The signs of the translational velocities (u, v, w).
- That the velocity corrections are reasonable and free of glitches caused by unreliable input data.

For more information, see *IRIS and RDA Utilities Guide*.

B.1.5.2 Using the Rays Utility to Test Velocity Correction

The **Rays** utility allows you to examine ingest data.

Also, because **RAW** products can be brought back from archive to create regenerated ingest files, you can check data recording.

You can use the **Rays** utility values to verify the velocity correction algorithm from an independent calculation.

B.1.6 In Situ Testing Suggestions

You must do the final testing at sea. This is an uncontrolled environment for testing, however, there is a useful total system test of the velocity correction. Fixed ground targets should have 0 corrected velocity as observed from the moving platform. Do not use buoys with radar reflectors. They are excellent targets, but they are generally not stationary.

Cruising offshore at ~10 km (6.2 mi) and observe fixed ground targets while collecting IRIS data. You may want to use a sector scan for this at low elevation. Verify the following:

- The real time color display should show ground targets at zero velocity.
- The **Rays** utility should show targets at zero velocity.
- **PPI** velocity products with a fine resolution velocity scale can be constructed to verify the velocity correction.

If the corrected velocity of ground targets is not 0, you must isolate the problem. If the antenna is properly stabilized, it is likely that both the INU (or Gyro system) and the antenna controller are functioning properly. Therefore, the error is most likely due to alignment:

- The INU or Gyro system vertical and the antenna vertical must be co-aligned.
- The INU or Gyro system azimuths must be co-aligned.
- The moment arms must be properly measured.

The INU or Gyro/Antenna alignment can be checked with the sun tracking feature of the **Antenna** utility. The moment arm measurement can be checked dynamically.

B.1.6.1 Dynamic Adjustment of Moment Arms

Measuring moment arms or extracting them from the ship's drawing can be difficult. The athwartship component of the moment arm is usually the easiest measure and, in most cases, the antenna is centered across the ship. This leaves the vertical and the for/aft moments.

You can check these moments dynamically by observing ground targets in special ship or antenna configurations.

1. Isolate the vertical component of the moment arm by scanning over the beam of the ship while looking at on-shore clutter targets.
   In this case, the velocity correction is dominated by the roll of the ship and the vertical moment arm. The athwartship position error of the antenna moment arm has the same effect. In this configuration, adjust the vertical moment arm in **Antsetup** until the observed velocity of ground targets is 0. The ship must be rolling during this test.

2. Point the antenna fore or aft to observe clutter targets.
   Adjust the fore/aft component of the moment arm to remove any remaining velocity error. The ship must be pitching during this test.

B.1.7 Velocity Correction Algorithm: INU Example

1. $V_R = V_R - V_{R\ ship}$
2. $V_{R\ ship} = x \cdot \dot{y}_A$
3. $y_A = \dot{y}_I + \dot{A}^{-1} \Delta y'_{AI}$

4. $y' = \begin{bmatrix} SpeedEast \\ SpeedNorth \\ Speed\,Up \end{bmatrix}$ $as\ measured\ by\ the\ INU$

5. $\Delta y_{AI} = antenna\,position \begin{bmatrix} forward \\ port \\ over \end{bmatrix}$ $the\ INU\ position$

6. $-1 = \begin{bmatrix} \dot{A}_{11}^{-1} & \dot{A}_{12}^{-1} & \dot{A}_{13}^{-1} \\ \dot{A}_{21}^{-1} & \dot{A}_{22}^{-1} & \dot{A}_{23}^{-1} \\ \dot{A}_{31}^{-1} & \dot{A}_{32}^{-1} & \dot{A}_{33}^{-1} \end{bmatrix}$

$$\dot{A}_{11}^{-1} = -\dot{P}\sin P\cos\theta - \dot{\theta}\cos P\sin\theta$$

$$\dot{A}_{21}^{-1} = +\dot{P}\sin P\sin\theta - \dot{\theta}\cos P\cos\theta$$

$$\dot{A}_{31}^{-1} = -\dot{P}\cos P$$

$$\dot{A}_{12}^{-1} = -\dot{R}\sin R\sin\theta + \dot{\theta}\cos R\cos\theta - \dot{R}\cos R\sin P\cos\theta - \dot{P}\sin R\cos P\cos\theta + \dot{\theta}\sin R\sin P\sin\theta$$

$$\dot{A}_{22}^{-1} = -\dot{R}\sin R\cos\theta - \dot{\theta}\cos R\sin\theta + \dot{R}\cos R\sin P\sin\theta + \dot{P}\sin R\cos P\sin\theta + \dot{\theta}\sin R\sin P\cos\theta$$

$$\dot{A}_{32}^{-1} = -\dot{R}\cos R\cos P + \dot{P}\sin R\sin P$$

$$\dot{A}_{13}^{-1} = +\dot{R}\cos R\sin\theta + \dot{\theta}\sin R\cos\theta - \dot{R}\sin R\sin P\cos\theta + \dot{P}\cos R\cos P\cos\theta - \dot{\theta}\cos R\sin P\sin\theta$$

$$\dot{A}_{23}^{-1} = +\dot{R}\cos R\cos\theta - \dot{\theta}\sin R\sin\theta + \dot{R}\sin R\sin P\sin\theta - \dot{P}\cos R\cos P\sin\theta - \dot{\theta}\cos R\sin P\cos\theta$$

$$\dot{A}_{33}^{-1} = -\dot{R}\sin R\cos P - \dot{P}\cos R\sin P$$

HEAD
    Heading measured CW from true north

θ
    HEAD - $\pi/2$

$\dot{\theta}$
    Rate of change of heading

P
    Pitch measured relative to horizon + for bow down

$\dot{P}$
    Rate of change of pitch

R
    Roll measured relative to the horizon + for port side down

$\dot{R}$
    Rate of change of roll

Appendix C. IRIS 3DView

C.1 IRIS 3DView Overview

IRIS 3DView software is a separately licensed software package that renders 3D visualizations of IRIS format radar data.

The radar data can be rendered using surface or volume rendering over a detailed 3D terrain map of the land and water underneath the weather echoes. The 3D visualization can also be rendered over a 2D radar product.

IRIS 3DView allows the rendered scene to be rotated, panned, and zoomed. You can add additional detail layers such as radar range rings, latitude-longitude grid, and landmark labels to the view for easy navigation.You can also store preferred view angles and display states for later use.

You can view the data files (Archive Mode) or configure IRIS 3DView to monitor a data file directory for viewing of the latest radar data (Live Mode). The input radar data may be data collected from a single weather radar, or from an IRIS composite of multiple radars.

You can export bitmap images of the 3D visualization, and run a background process that automatically renders bitmap images of the incoming radar data.
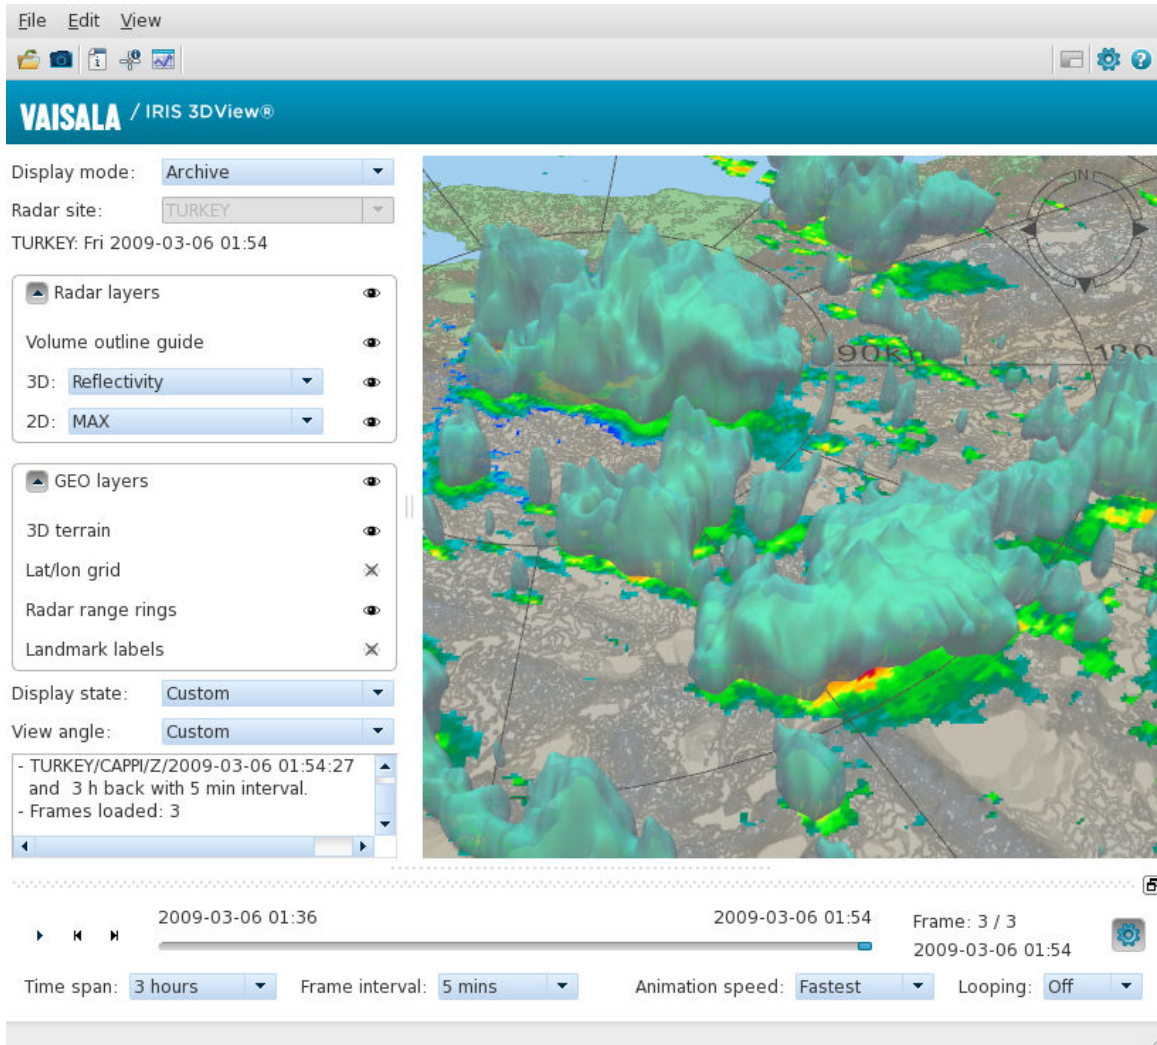
Figure 105   IRIS 3DView User Interface

C.1.1 IRIS 3DView in System Architecture

The IRIS 3DView package fits in the IRIS system architecture as shown in the following figure.

Generally, there is an IRIS/ Radar package running the IRIS software at a remote radar site (or perhaps there are multiple remote radar sites).

Raw radar base data are fed from the radar sites to a central server computer running IRIS/ Analysis software generating many 2D radar products. These products may be viewed directly on this central server computer, or re-distributed to other client computers running the IRIS/Display software package. For 3D functions, the central server computer outputs radar products over the network to the IRIS 3DView workstation. The IRIS 3DView computer receives these data and displays it in the 3D rendering window.

At the 3D rendering window, the operator may manipulate the data by rotating, panning, zooming, and so on.

During manipulation, the images are viewed in true 3D form. Any 3D image viewed on the IRIS 3DView computer can then be sent back to the central server computer for viewing and distribution to other client computers on the network running IRIS/ Display.

The 3D functions (rotation, etc.) are only available at the IRIS 3DView workstation. However, 2D versions of the 3D renderings can be made available to any IRIS workstation on the network.
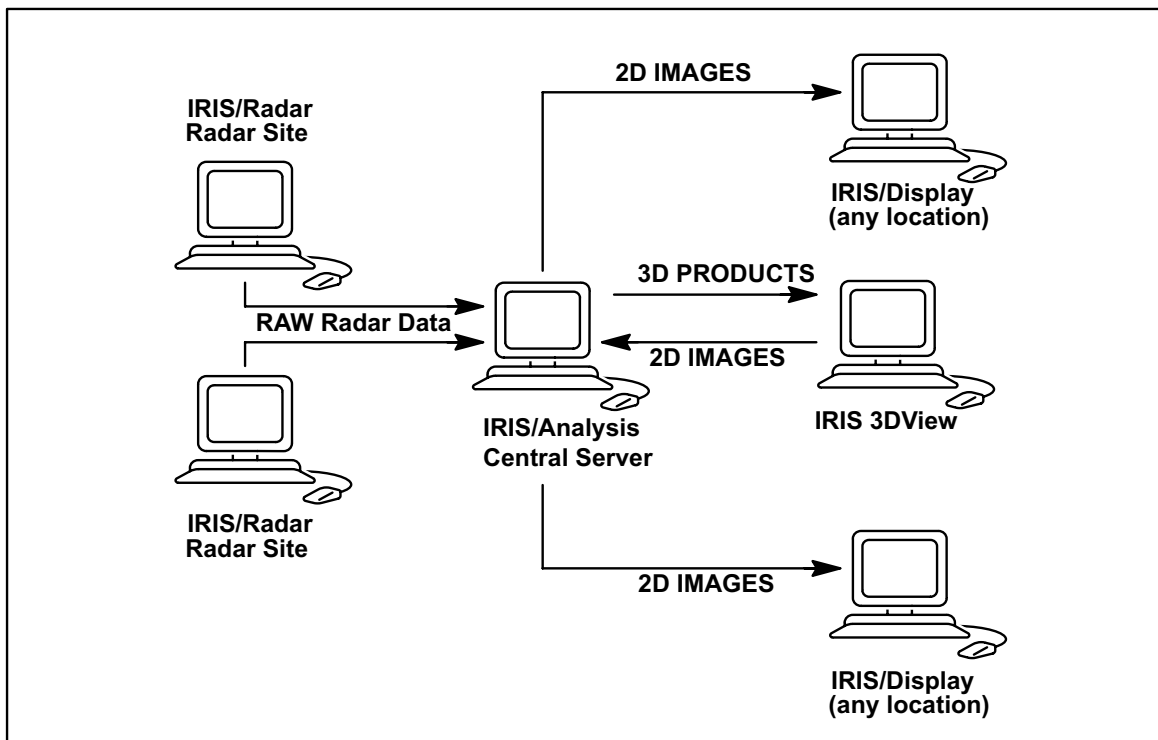
Figure 106   IRIS 3DView Data Flow

C.1.2 IRIS 3DView System Requirements

IRIS 3DView requires a Linux operating system. The installation files are delivered in RPM (Red Hat Package Management) format.

IRIS 3DView runs on standard PC hardware. IRIS 3DView must be run on a dedicated computer, since it requires full use of the processor, memory, and graphics subsystems.

The minimum requirements for the hardware are as follows:

- Processor: Intel Core 2 Duo (3.0 GHz, 6 MB L2 cache, 1333 MHz FSB)
- Memory: 2 GB DDR2 800 MHz
- Graphics card: nVidia GeForce 9600 or equivalent

C.2 Installing IRIS 3DView

It is not necessary for IRIS software to be installed on the system to use 3DView.

If you chose to have the IRIS software on the same system, it is recommended that you install the IRIS software before installing the IRIS 3D View.

IRIS 3DView software is a separately licensed software package that renders 3D visualizations of IRIS format radar data.

You can install the IRIS 3DView either on an IRIS Analysis machine, or a machine without IRIS.

1. Obtain the following RPMs:
   - *libogg-devel-1.1.4-2.1.el6.i686.rpm*
   - *libtheora-devel-1.1.0-2.el6.i686.rpm*
   - *qtwebkit-2.1.1-1.el6.i686.rpm*
   - *qt-mobility-1.1.3-2.el6.i686.rpm*
   - *gl2ps-1.3.5-1.el6.i686.rpm*
   - *gl2ps-devel-1.3.5-1.el6.i686.rpm*
   - *compat-expat1-1.95.8-8.el6.i686.rpm*
   - *expat-2.0.1-9.1.el6.i686.rpm*
   - *expat-devel-2.0.1-9.1.el6.i686.rpm*
   - *xorg-x11-apps-7.4-10.el6.i686.rpm*
   - *xorg-x11-server-common.1.7.7-29.el6.i686.rpm*
   - *xorg-x11-server-Xvfb-1.7.7-29.el6.i686.rpm*
   - *vtk-5.8.0-6.el6.i686.rpm*
   - *vtk-devel-5.8.0-6.el6.i686.rpm*
   - *vtk-java-5.8.0-6.el6.i686.rpm*
   - *vtk-python-5.8.0-6.el6.i686.rpm*
   - *vtk-qt-5.8.0-6.el6.i686.rpm*
   - *vtk-tcl-5.8.0-6.el6.i686.rpm*
   - *acroread-9.4.0-1.el6.i686.rpm*
   - *acroread-plugin-9.4.0-1.el6.i686.rpm*

2. Install the RPMs in the following order:
   a. **#rpm -Uvh** `libogg-devel-1.1.4-2.1.el6.i686.rpm`
   b. **#rpm -Uvh** `libtheora-devel-1.1.0-2.el6.i686.rpm`
   c. **#rpm -Uvh** `qtwebkit-2.1.1-1.el6.i686.rpm`
   d. **#rpm -Uvh** `qt-mobility-1.1.3-2.el6.i686.rpm`
   e. **#rpm -Uvh** `gl2ps-1.3*`
   f. **#rpm -Uvh** `compat-expat1-1.95.8-8.el6.i686.rpm`
   g. **#rpm -Uvh** `expat-2*`
   h. **#rpm -Uvh** `xorg-x11*`
   i. **#rpm -Uvh** `vtk*`
   j. **#rpm -Uvh** `acroread-9*`

3. Log in as **root**.

4. The IRIS 3D View application is in the rpm package called *iris3dview-1.0.2-1.i686.rpm*. Install it using the rpm installation tool

```
# rpm -Uvh iris3dview-1.0.2-1.i686.rpm
```

> **i** It is recommended that you do not run the IRIS 3D View application while logged in as a **root** user.
> If you have IRIS software installed on your system, log in as **radarop** user.
> If IRIS software is not installed, create a user on your system.

5. In the **GNOME** window, create a user by selecting **System > Administration > Users and Groups** and follow the online instructions.

6. Log in as normal user and check *.config/Vaisala* directory for the following files:
   - *IRIS3D.conf*
   - *IRIS3DColors.conf*
   - *IRIS3DLandmarks.conf*
   - *IRIS3DVisStyles.conf*

7. Do one of the following:

   a. If you have the IRIS software installed, create the following directories:

   ```
   /usr/iris_data/data/geo      # for geo data
   /usr/iris_data/data/2d       # for 2d data
   /usr/iris_data/data/3d       # for 3d data
   /usr/iris_data/data/background_output
   # for background images processing results
   ```

   The permissions and ownership for these directories are:

   ```
   drwxrwsr-x
   2 operator users 4096 Dec 11 13:14 2d drwxrwsr-x
   2 operator users 4096 Dec 11 13:15 3d drwxrwsrwx
   2 operator users 4096 Dec 11 13:15 background_output drwxrwsr-x
   2 operator users 4096 Dec 11 13:14 geo
   ```

   b. If you do not have the IRIS software installed, you can put your data directory anywhere you want, but you must set the directory permissions to read/write access.

8. Run **IRIS 3DView** visualization:

```
/usr/bin/iris3dview &
```

When you first start **IRIS 3DView**, it runs a license wizard. Follow the online instructions.

9. If your system has an IRIS or RDA system already installed, the **IRIS 3D View License** dialog box displays the **Sitename** for your system.
In that case, leave it unchanged.
If not, enter a unique **Sitename** (16 characters or less) and select **Next**.
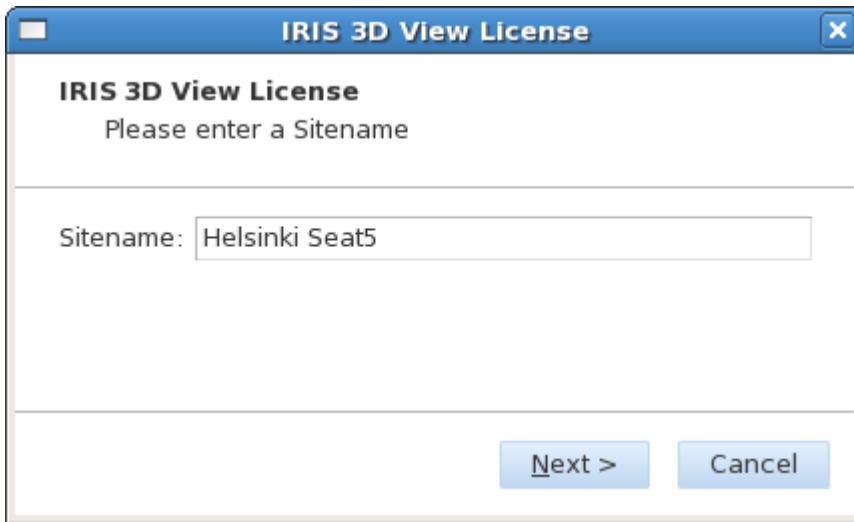
The **IRIS 3D View License** dialog displays the site code and check code.
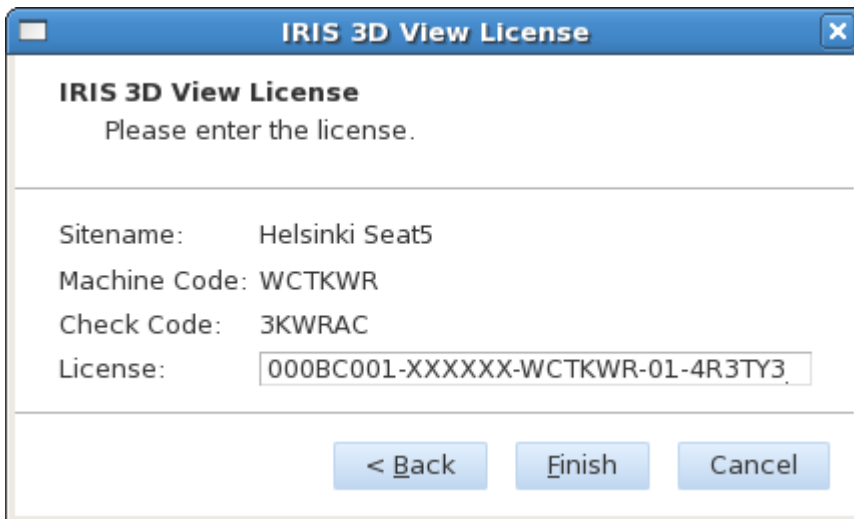
10. Send the information shown in previous figure to Vaisala and request a license string.

11. Enter the license string in the **License** field and click **Finish**.

12. Once the **IRIS 3D View** window opens, click the setup icon in the upper right hand corner to open the **Settings Window** dialog box.
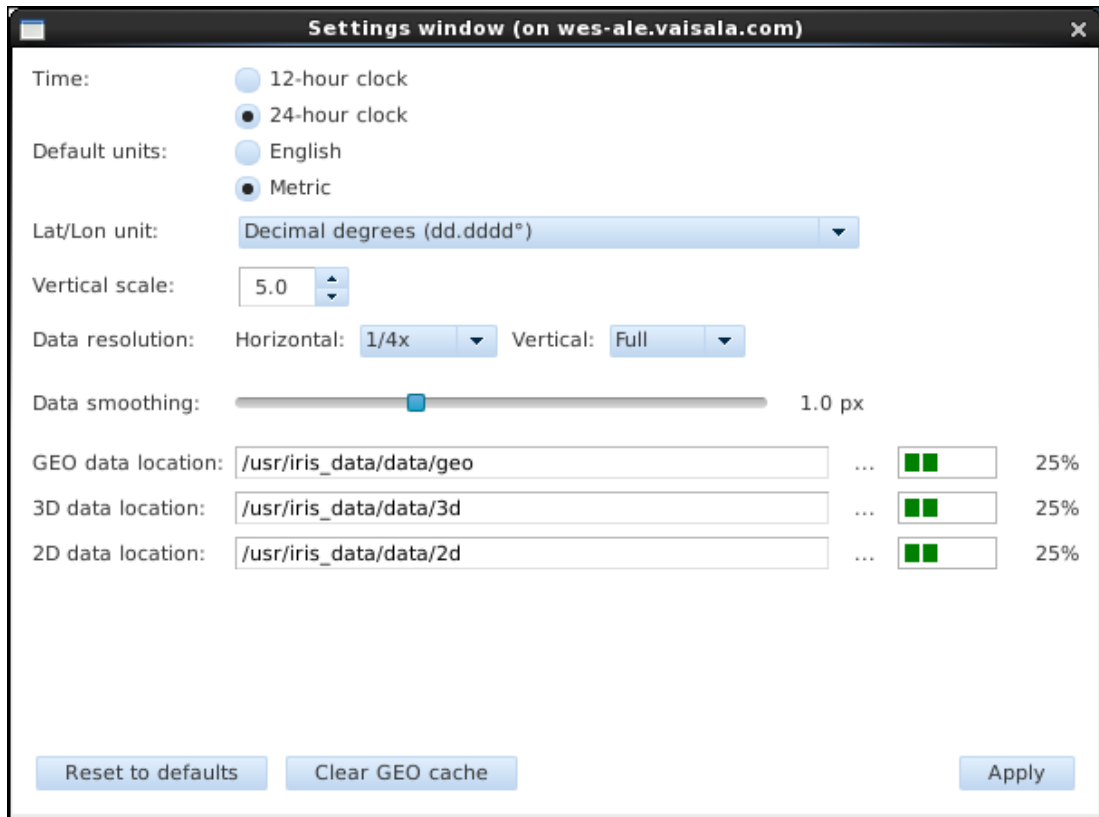
Configure the data directory to point to the directories created earlier, according to your data type.

13. Configure background processing.
    Background processing is run by the `iris3d` user; however, the user who runs the
    `iris3d` visualization must copy the configuration files to the `.config` directory of
    the `iris3d` user:

    a. Set permissions for everybody to be able to read and write by running:

    ```
    #chmod -R 777 /home/iris3d/ as root user
    ```

    b. From the visualization window, select **File > Export tool**.

    c. From the export window, select the **Background** output tab and follow the online
       instructions to finish the setup.
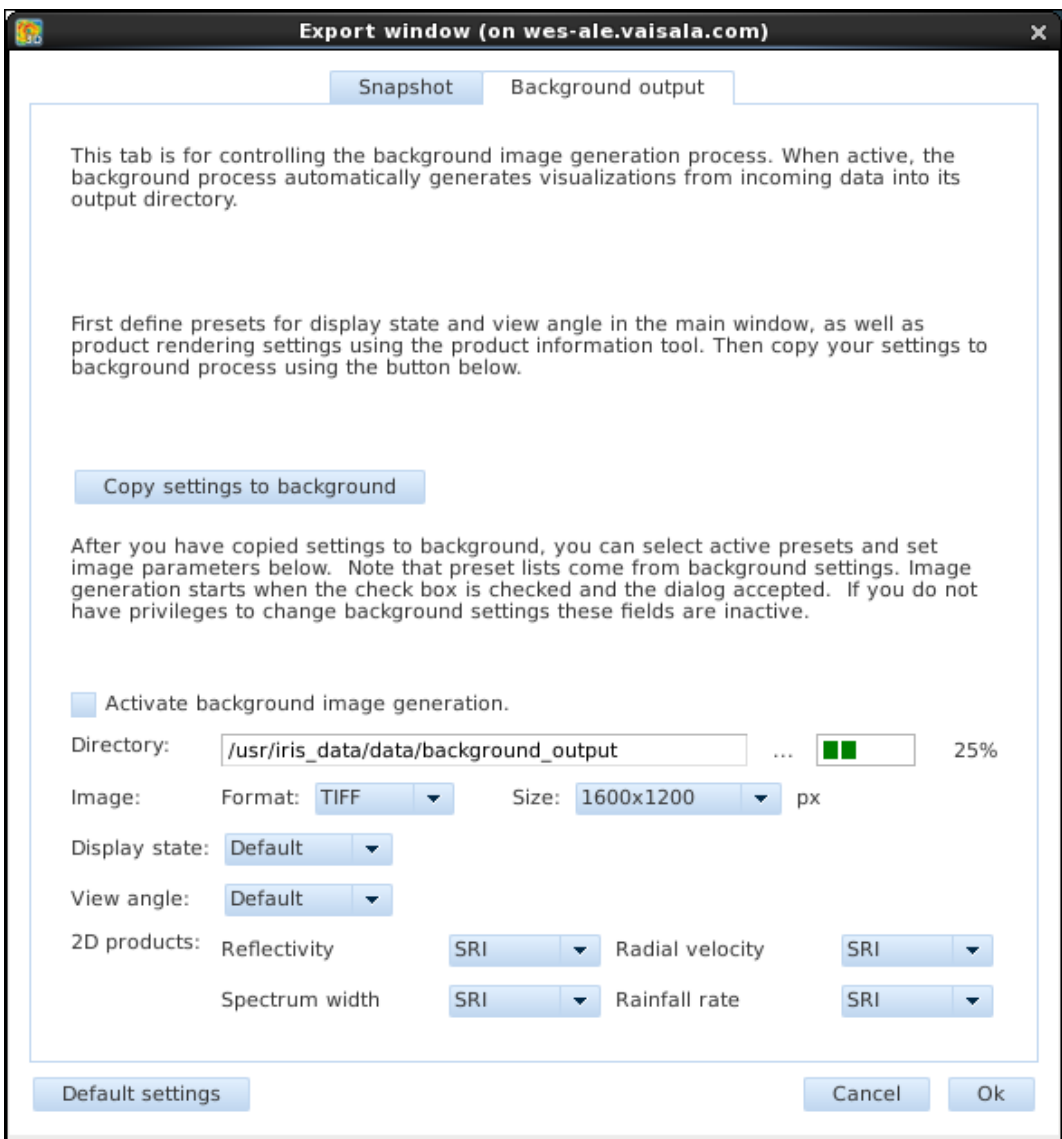
After installing the IRIS 3D View, the background processing daemon starts automatically.

14. Check the daemon by running the `ps -ef` command:

```
ps -ef | grep iris
```

and you should get the printout similar to:

```
iris3d   3432    1  0 13:24 ?       00:00:00 dbus -launch -- autolaunch
841d3b9bca80c846219d29fc00000022 --binary-syntax --close- stderr
iris3d   3433    1  0 13:24 ?       00:00:00 /bin/ dbus - daemon --fork --
print- pid 5 --print-address 7 --session
iris3d   3435    1  0 13:24 ?       00:00:00 /usr/bin/ iris3dviewd --
loglevel =2 --daemon
```

If you need to restart the daemon, change to the root user by typing:

```
su root
```

You are prompted for the password.
Stop and start the daemon by typing the following:

```
#/sbin/service iris3dviewd stop
#/sbin/service iris3dviewd start
```

15. Test background processing:

   a. Copy 3D data to the 3D directory that you configured.

   b. Wait for a second and check the output file in output directory that you configured.

16. If needed, modify the logging level.

   a. To check the logging, go to */home/iris3d/* and open the *iris3dviewd.log* file.

   b. To view real-time logging messages while you are running background processing, run `tail -f iris3dviewd.log`.

   c. To set logging level to more details, like level 6, editing the */etc/sysconfig/ iris3dviewd* and change `loglevel=2` to `OPTIONS="-- loglevel=6 -- daemon"`.

   d. Restart the processing daemon.

C.3 Getting Started with IRIS 3DView
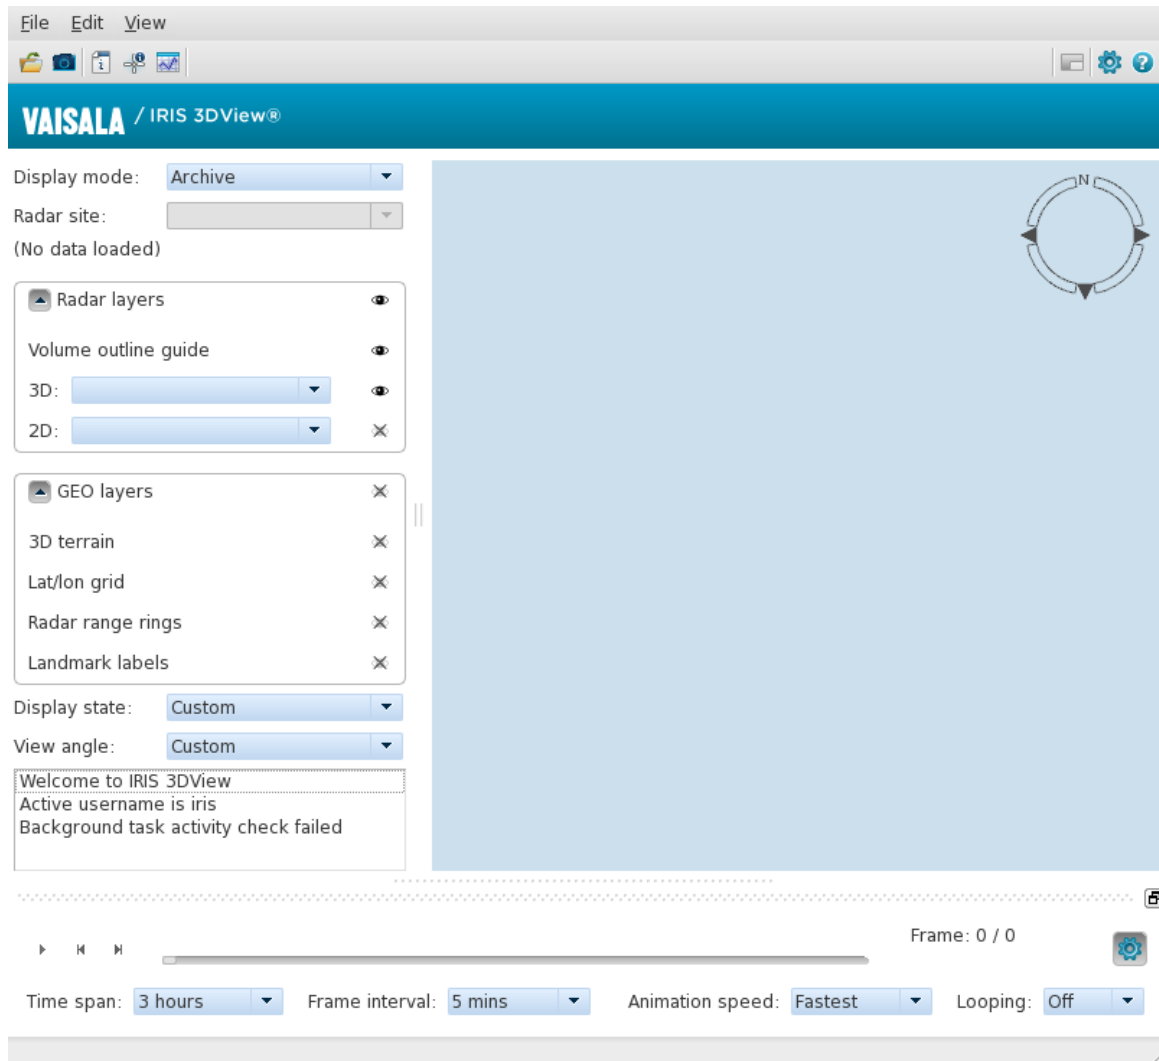
Figure 107   IRIS 3DView Starting View

▶ 1. On your IRIS 3DView workstation, select the IRIS 3DView icon.
    Alternatively, in a shell session, type `iris3dview`
    The software starts and the main window is shown with controls for the view on the left.
    The starting view is a direct downward view.
    The volume outline guide (dark lines) are drawn around the rendered area.

2. Move the mouse pointer over the compass ring to show the current view heading and zoom level.

3. Use your mouse to control the 3D view:
   - Hold the left mouse button and drag the mouse to move the camera position about the center of the picture.
     The camera remains aiming at the center, so this effectively rotates the data.
     Drag the mouse horizontally to rotate the camera about the center.
     Drag the mouse vertically to move the camera vertically about the center.
   - Hold the right mouse button and drag the mouse to move the camera location.
     The camera pointing angle remains the same so this pans the data. Holding the left button and shift key does the same.
   - Hold the middle mouse button and drag the mouse (or rotate the mouse wheel) to move the camera closer or farther from the center. This effectively zooms the data.

4. To save your own view angle presets, use the view angle preset manager.
   See C.6 Managing the IRIS 3D Display (page 336).

5. To return to the starting view at any time, select the default angle preset.

C.4 Opening Data Files in IRIS 3DView

1. Make sure the data file locations are configured in IRIS 3D **Settings**.
   See C.12 Managing IRIS 3DView Settings (page 345).

2. To open a data file, select **File > Open 3D data**.
   This opens the 3D data selector where you can select the data file.

3. To open the file in **Live** mode, select **Display mode > Live**.
   In **Live** mode, IRIS 3DView opens the latest data file from the radar site selected in the **Radar site** drop-down list.
   See C.6 Managing the IRIS 3D Display (page 336).

C.5 Rendering Modes

Rendering mode determines the visualization approach that IRIS 3DView takes when creating the 3D model of the radar data.

You choose the rendering mode and the mode-specific options in the **Settings** window. For rendering, you can choose from surface rendering and volume rendering.

IRIS 3DView uses a full spectrum color scale when rendering the radar data. Blues represent the weakest echoes, building up to greens, yellows, oranges, and finally reds. The color scale is NOT absolute, it is relative. The reflectivity legend on the product information tool window shows the current relationship of colors and reflectivity in decibels (dBZ).

C.5.1 Surface Rendering

Surface rendering mode uses what in 3D terminology are called iso-surfaces. In other words, for example, if the threshold of the first iso-surface is set at 30 dBZ, all of the echo that is less than 30 dBZ is removed from the picture and what you see is a shape that represents the outer boundary of the 30 dBZ echo.

IRIS 3DView allows up to 4 of these iso-surfaces to be drawn in the same 3D visualization. The threshold value, color and opacity of each surface is user configurable. This allows a flexible representation of the radar data.

Proper opacity control of the surfaces is important for good visualization. The more opaque a surface is (higher opacity value), the more difficult it is to see through. The less opaque, the more transparent the surface becomes and it is easier to see through it.

Figure 108   Example of Surface Rendering

C.5.2 Volume Rendering

In volume rendering mode, IRIS 3DView renders a 3D visualization of the data that is between the minimum and maximum values. All other data outside this range is filtered out and not displayed. Instead of using distinct surfaces, the area is rendered so that it gradually changes from the color and opacity at the minimum value to the color and opacity at the maximum value.

The minimum and maximum values, as well as the color and opacity at the ends of this range, are user configurable.

Figure 109   Example of Volume Rendering

C.6 Managing the IRIS 3D Display
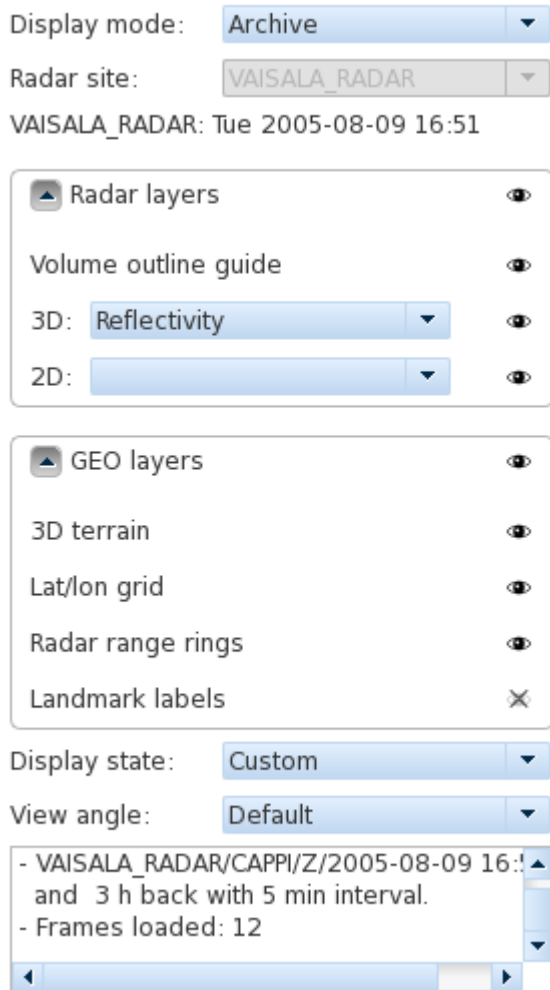
Figure 110   IRIS 3D Display

Use the options on the left of the main window to define what is shown in the 3D view.

> Drag the handle on the right side of the control area to expand the control area.
> Toggle the "eye" icons to manage the visibility of the controlled visual element or layer.

1. In **Display Mode**, select the operating mode:
   - **Archive** mode is a manual mode where you can loads radar data files and creates visualizations without any automatic functions taking place.
   - In **Live** mode the IRIS 3DView software continuously monitors the input directory. When a new radar data file comes in from the central server computer, IRIS 3DView automatically loads it.
     If the background process has been configured and enabled, it renders 2D images of the radar data in the background, using preset view angles and display states. The background process requires some setup (that is, setting of a pre-defined view point, and similar), but no interaction is required at runtime.

2. If you are in **Live** mode, in **Radar Site**, select the monitored radar site.
   The radar site and timestamp of the currently loaded data file are shown.

3. In **Radar Layers**, toggle the **Volume outline guide** on or off (the dark lines around the rendered area), and select the 3D and/or 2D radar products that are rendered.

4. In **GEO Layers**, select which if the following geographical information layers are shown:
   - 3D terrain
   - Latitude-longitude grid
   - Radar range rings
   - Landmark labels

5. In **Display State**, do one of the following:
   - Select a display preset from the drop-down list.
   - Select **Manage presets** to open the preset manager window.
     In the preset manager, you can store the currentdisplay state options as a preset, and manage existing presets.

6. In **View Angle**, do one of the following:
   - Select a view angle preset from the drop-down list.
   - Select **Manage presets** to open the preset manager window.
     In the preset manager, you can store the current view angle as a preset, and manage existing presets.

The text window at the bottom of the controls shows a running log of the configuration changes.

C.7 Using the IRIS 3DView Export Tool

Use the export tool to save snapshots of 3D views as bitmap images.
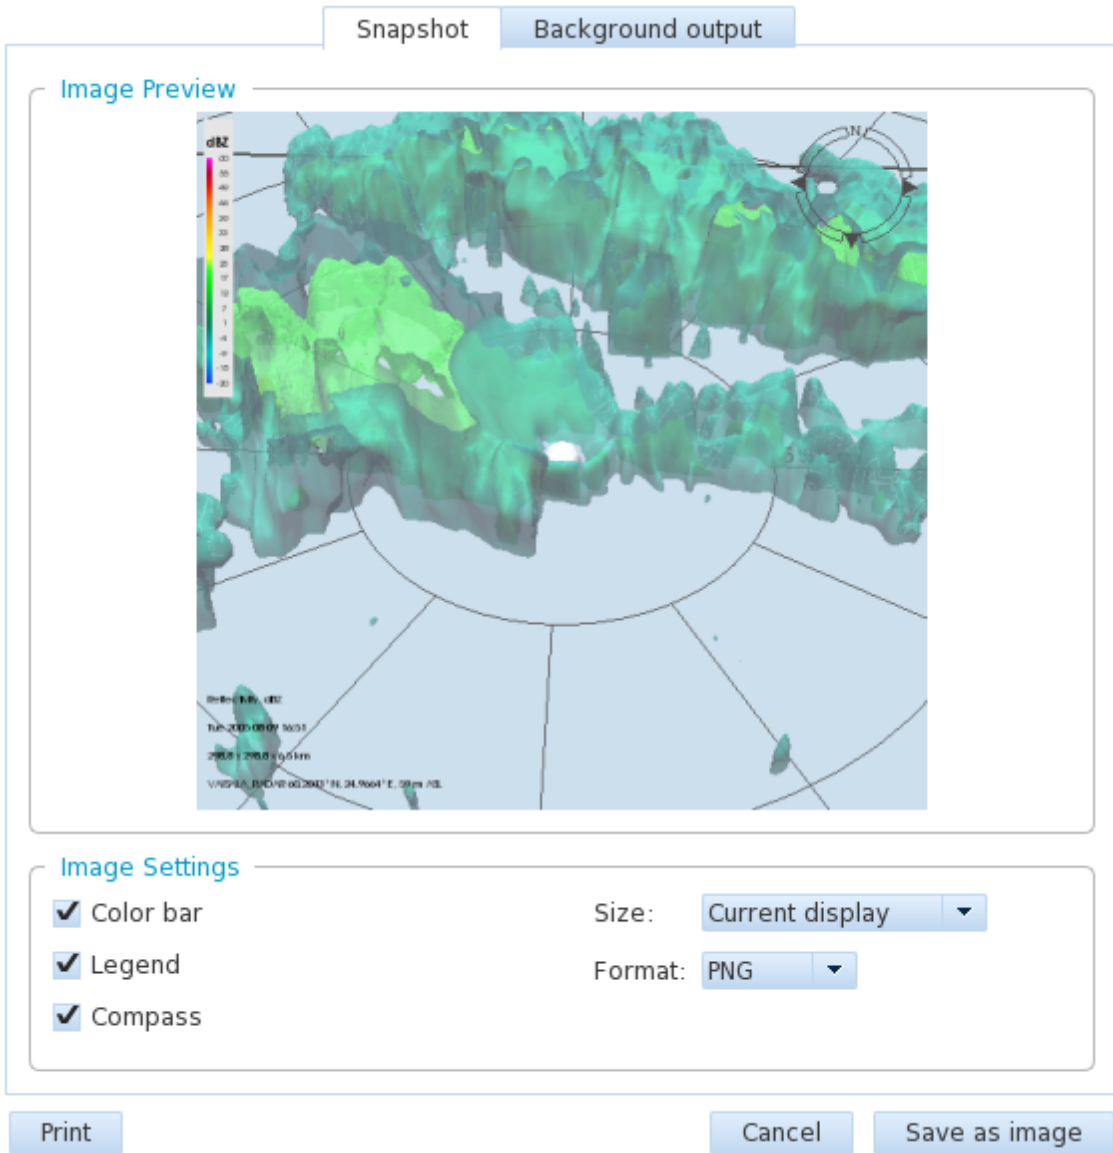
Figure 111   Export Tool

1. Select **File >** 📷.
    The export tool window opens with the currently rendered 3D view shown with a reflectivity legend.

2. To save the current image as a file, select **Save as image**.

3. In the **Background output** tab, enable and configure the background rendering process that outputs 2D images using pre-defined options:
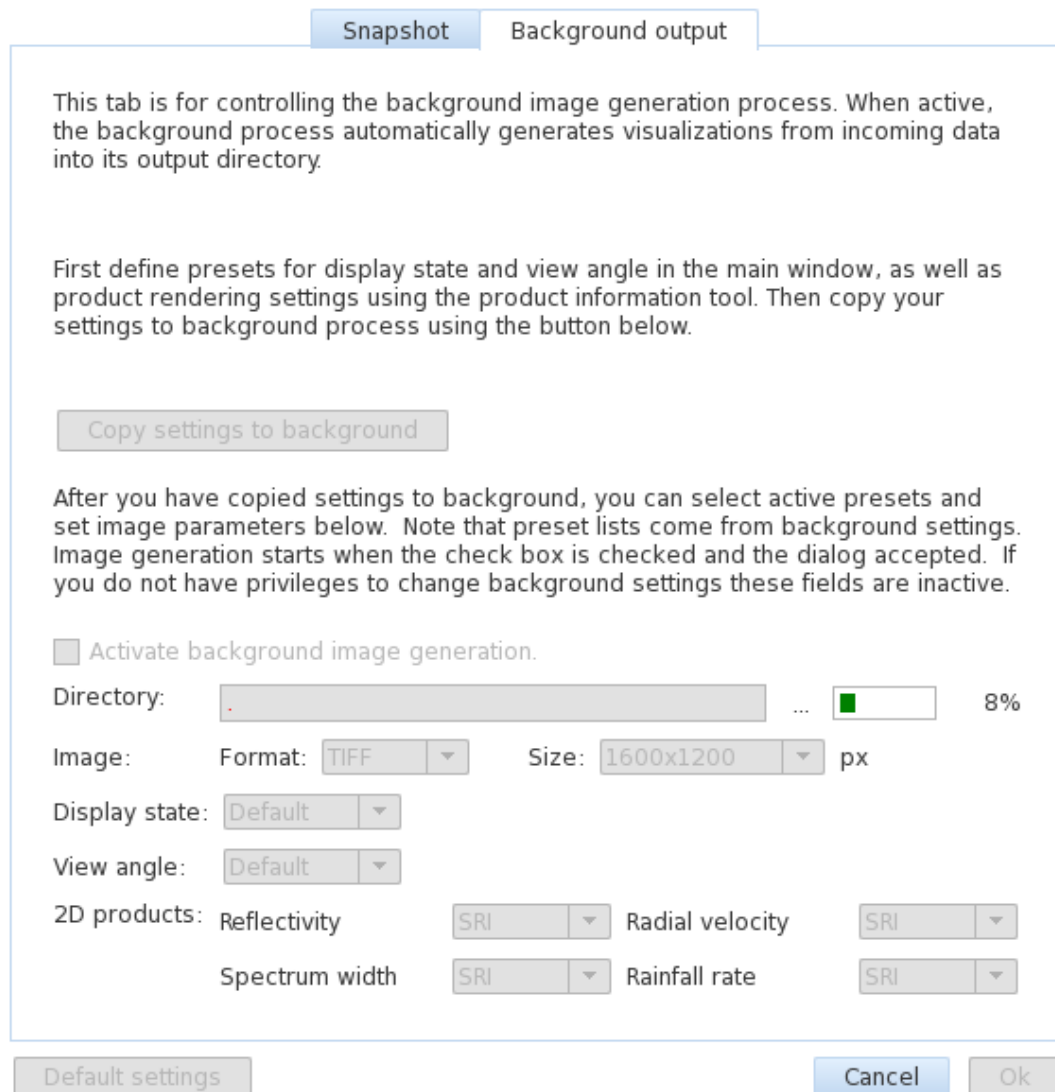
Figure 112   IRIS 3DView Background Output Configuration Window

a. Create and save a preset for the **Display state** used in background rendering.

b. Create and save a preset for the desired **View angle**.

c. Select **Copy settings to background** to make the presets available for selection.

d. Set the desired output generation options, including output directory, file format, and image size.

e. Select **Activate background image generation**.

f. Select **OK** to exit to the main view.

C.8 Using the Product Information Tool

Figure 113   Product Information Tool

▶

1. Select **File >** .
   The product information window opens with the following information about the
   currently selected 3D data product:
   - **Data**:
     Data type being displayed
   - **Time**:
     Timestamp of the data file
   - **Volume**:
     Size of the radar data volume
   - **Radar**:
     Name and location of the radar that produced the data file
   - Color legend for the 3D view

2. To define how the data product is rendered in the 3D view, use the following options:
   - **Color map**:
     Allows selection from a short list of pre-configured color scales.
   - **Rendering mode**:
     Rendering mode of the 3D visualization. Options are: surface rendering and volume
     rendering.
   - **Isosurface parameters**:
     These options are visible if surface rendering has been selected as the rendering
     mode. Using these options you can define the following options for surface
     rendering:
     - Number of isosurfaces (1 ... 4)
     - Threshold value for each isosurface
     - Opacity of each isosurface
     - Color of each isosurface

3. If you have selected volume rendering as the rendering mode, you can define the
   following **Volume parameters** for volume rendering
   - Low and high value of the rendered range
   - Opacity at the low and high end of the rendered range
   - Color at the low and high end of the rendered range

4. Do one of the following:
   - Select **Apply** to save your changes.
   - Select **Revert to previous** to restore the previous settings.
   - Select **Reset to defaults** to select the surface rendering mode with the default
     isosurface settings.

5. To minimize the window, select the small arrow at the bottom right corner.
   In the minimized form, the product information tool shows only the color legend.

C.9 Using the IRIS 3DView Cursor Tool

The cursor tool displays information about the current home location, and the current cursor
position in relation to the home position.
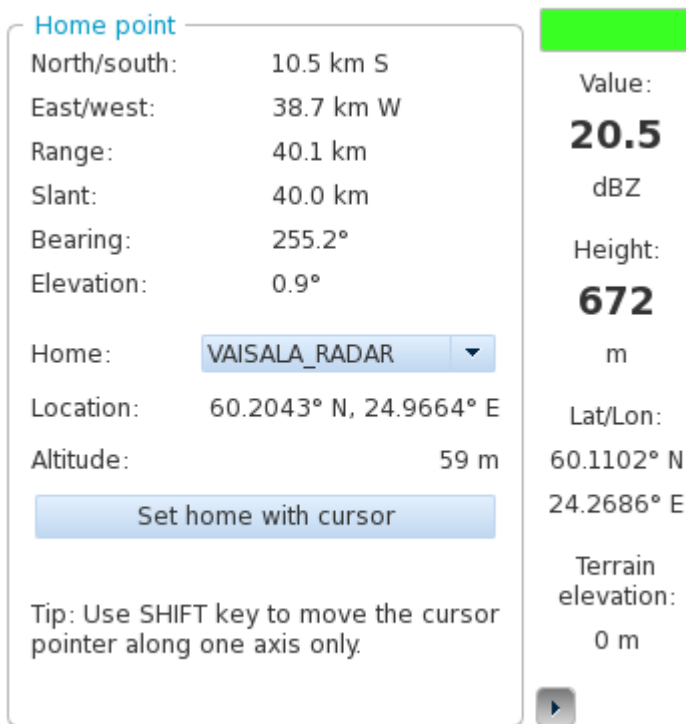
Figure 114   Cursor Tool

1. Select **File >** .
   The cursor tool window opens and a round cursor with yellow axis lines appears on the
   3D view area.

2. Use the mouse to drag the cursor around the view area.
   Position the mouse pointer above the cursor or one of the axis lines, hold down the left
   mouse button, and move the mouse.
   As the cursor moves, the cursor tool window is updated with the location and value
   information at the cursor position.

> To move the cursor only on horizontal or vertical axis, hold down the shift key when
> moving the cursor.

> You can also change the current home position by selecting a radar site from the
> drop-down list, or selecting **Set home with cursor**.

3. To minimize the window, select the small arrow at the bottom right corner.
   In the minimized form, the cursor tool only shows the information for the current cursor
   position.

C.10 Adjusting IRIS 3DView Screen Settings

▶

1. To switch between a fullscreen mode and the normal windowed mode, select ⬚.

2. To switch between bright and dark color schemes, select ⬚.

C.11 Playing Back Data

Figure 115   IRIS 3DView Animation Toolbar

The animation toolbar controls the playback of data frames.

The top of the toolbar shows the timestamps of the first and last frames of the currently loaded data file, the number of the current frame, and the total number of frames created from the data file. The number of frames is affected by the amount of data in the file, and the animation settings that are active on the toolbar, such as **Time span** and **Frame interval**.

The slider selects the currently rendered frame. You can drag the selector using the mouse. Use the playback controls on the left of the slider to play the animation from the current frame to the end, or to jump to the next and previous frames. You can separate the animation toolbar from the main window by pressing the icon on the top right corner of the animation toolbar.

1. To control animation playback settings, select the settings icon ⚙ on the animation toolbar.

    Time Span
    > Determines the maximum time selectable using the animation slider. The options are 3, 6, 12, or 24 hours back in time.

    Frame Interval
    > Determines the minimum frame interval. If there are more frequent data files on disk, then some are skipped to reach this minimum.
    > The larger the frame interval, the faster the animation proceeds through the time span.
    > The options are 5, 15, 30, or 60 minutes.

    Animation Speed
    > Sets an additional delay between displayed frames. Does not affect the rendering time. Options are:
    > - **Fastest** (no delay)
    > - 1 second
    > - 3 seconds
    > - 5 seconds
    > - 10 seconds

    Loop mode
    > Select to have the animation **Play once** or **Loop** until stopped.

> ℹ The typical bottleneck for animation playback is CPU power, that is, the time it takes for IRIS 3DView to calculate the 3D model of the frame.
> If the animation playback speed is unsatisfactory, set the **Animation speed** to **Fastest**, render it in a smaller window, and consider lowering the data resolution in the **Settings** window.

C.12 Managing IRIS 3DView Settings

Figure 116   IRIS 3DView Settings Window

▶

1. In the top menu bar, select settings ⚙. The Settings window opens containing general settings for the IRIS 3DView software.

2. Adjust the settings as needed.

Time

Select **12-hour clock** or **24-hour clock**.

Default Units

Select **English** or **Metric** units.

Lat/Lon unit

Select the units for latitude and longitude. The available options are:
- **Decimal degrees (dd.dddd°)**
- **Degrees, decimal minutes (DD° mm.mmm')**
- **Degrees, minutes, decimal seconds (DD° MM' ss.s")**

Vertical scale

Set the vertical scale of the 3D rendering. A higher value will produce higher peaks, a low value will produce a flater rendering.

Data Resolution

Set the horizontal and vertical resolution that is used for rendering. Lowering the resolution will speed up the generation of the 3D model, but will produce a less detailed rendering.

Options for horizontal resolution:
- **Full**
- **Half**
- **1/3x**
- **1/4x**

Options for vertical resolution:
- **Full**
- **Half**

Data Smoothing

Increase data smoothing to remove some fine detail from rendered object, creating a smoother model. The smoothing value can be selected using a slider between 0 and 3 data pixels, at increments of 0.1 pixel.

Data Location Directories

Set the correct directories for the data files:
- **GEO data location:**
  The directory where IRIS 3DView looks for 3D terrain data.
- **3D data location:**
  The directory that IRIS 3DView scans for 3D radar data.
- **2D data location:**
  The directory that IRIS 3DView scans for 2D radar data.

3. Select **Apply**.

Appendix D. IRIS TDWR Features

D.1 Items Provided with IRIS/TDWR Option

The optional IRIS/SHEAR package includes features for terminal Doppler weather radar applications:

- **SHEAR** Product
  The **WARN** product run on a radial **SHEAR** product is used for microburst detection.
  See 5.6 SHEAR: Wind Shear (page 183).
- **SLINE** Product
  Detects wind shifts associated with gust fronts and cold fronts. Can use radial, azimuthal or combined radial and azimuthal shear.
  See 5.7 SLINE: Shear Line (page 195).
- Ribbon Display
  Ribbon display terminals (RBDT's) are large format text-only displays that are used to display and signal alerts to controllers. IRIS supports US FAA-style Dale RBDT's which are serial line hardware devices and Virtual RBDT's (vribbons) which are X-Window displays.
  See D.7 Ribbon Displays (page 355).
- TDWR/LLWAS Integrator
  Uses the US FAA algorithm to integrate information from IRIS shear warning products and LLWAS (low-level wind shear alert system) to generate TDWR products for output to standard US FAA-style ribbon displays. Note that an LLWAS system is not required.
  See D.8 TDWR/LLWAS Integrator and Runways Utility (page 366).
- Runway Utility
  Graphical utility for selecting the active runways that are in use. The information is used by the TDWR/LLWAS Integrator to select for which runway corridors alerts are issued to the RBDT's.
  See D.8.3 Runways Utility (page 373).

These features, combined with all the standard features of IRIS, provide controllers with timely, concise information concerning wind shear and other weather phenomena that can affect air terminal operations.

The TDWR used by the US FAA is based on technology circa 1985. Since this time, computer speeds have increased dramatically. There have also been changes, in algorithms, networking standards and user interface. The goal of the IRIS/TDWR is to take the best elements of the TDWR designs and implement them in the context of today's technology.

Vaisala does not warrant that the IRIS/TDWR option will function in accordance with US FAA Specifications.

D.2 IRIS/TDWR Terminology

Below is a summary of US FAA TDWR terminology and how it corresponds to IRIS terminology:

Arena

In IRIS this is called a protected area. For example runway `33LA_1MF` signifies a box (typically 1 miles square) that is centered 1 nautical mile from the approaching end (south end) of runway 33 (nominally approached on a heading of 330°).

Base Data

The basic `dBZ`, `V`, and `W` data. In IRIS, the compressed version is the **RAW** product. The un-compressed version is called Ingest Data.

Corridor

Each runway has two corridors, an approach corridor and a departure corridor.

Maintenance Display Terminal (MDT)

Used for control and monitoring of the TDWR system. Since IRIS is a networked client-server system, the concept of a dedicated "terminal" for this function does not really apply, since any workstation on the network can provide this function. For management purposes, it is sometimes useful to have an IRIS/Display workstation that is dedicated to the MDT task.

Microburst Alert (MBA)

Alert of a strong microburst. In IRIS, these are based on the **WARN** product run on a radial **SHEAR** product.

Low Level Wind Shear Alert System (LLWAS)

Based on anemometers that are positioned around the air terminal. LLWAS is not available at all sites. Where it is available, it provides backup to the TDWR in cases where the radar coverage is obscured, or there are no scattering targets.

Radar Data Acquisition System (RDA)

The data collection system. In IRIS, this consists of the signal processor, radar control processor and a workstation running IRIS/Radar software.

Radar Products Generator (RPG)

In TDWR, this is a separate computer that takes in Base Data from the RDA and generates products.

In IRIS, this is a workstation running IRIS/ Analysis receiving data from another workstation running IRIS/ Radar. In IRIS, the RDA and RPG functions can be combined.

Ribbon Display (RBDT)

A text only display intended for controllers. It connects to an IRIS Workstation through a serial line.

Situation Display (SD)

A simplified weather display showing oval shapes to depict microbursts and current and forecast wind shear line positions. In IRIS, this is accomplished by the simultaneous display of **WARN** and **SLINE** products.

Wind Shear Alert (WSA)

Alert of a shear line or gust front based on the **SLINE** (shear line) product, or a weak microburst based on the **WARN** product run on radial shear.

D.3 IRIS/TDWR Hardware Configuration Example

Vaisala supports single and dual–redundant hardware configurations. Because of the network communications features in IRIS there are many possible hardware configurations that allow different tasks to be performed by various workstations.

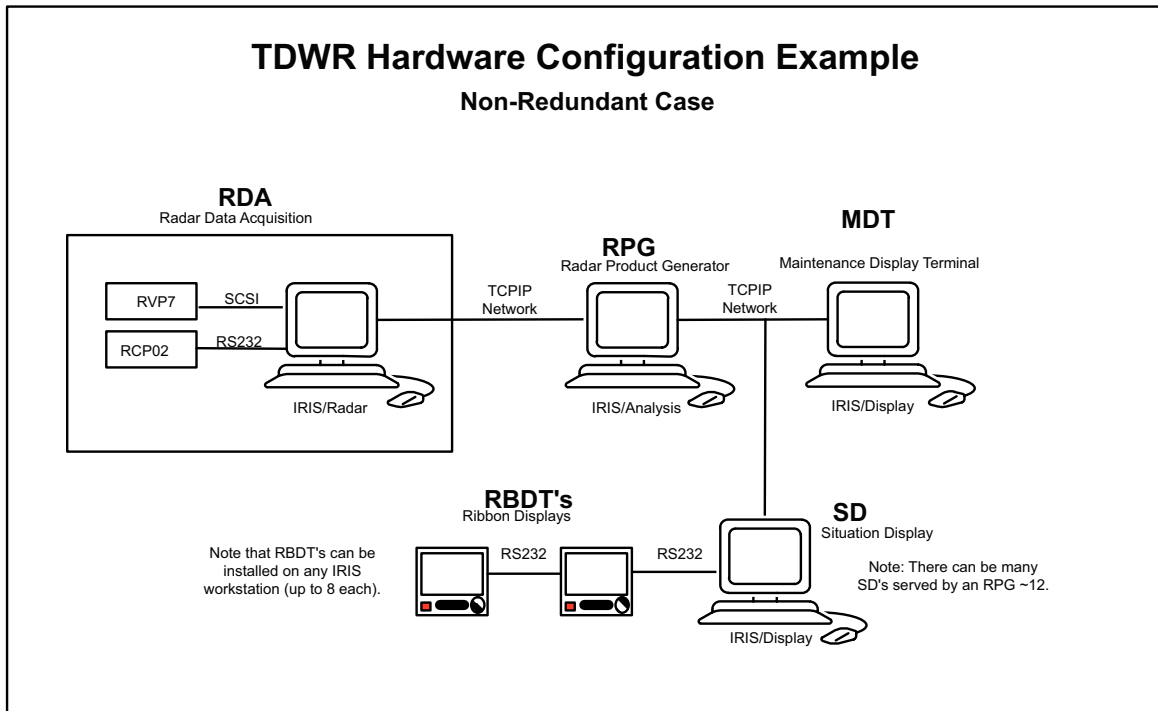The following example for a non-redundant system.

Figure 117   IRIS/TDWR Hardware Configuration Example

With the IRIS architecture, the RPG and MDT functions could be combined into the RDA workstation. Reducing the amount of hardware increases the system reliability and reduces maintenance. The RDA workstation also has all the features of an SD.

For redundant systems, there are two RDA's and two RPG's. The RDA's have dual network interfaces and the active one outputs base data to both RPG's. This means that all the products are made twice. Each SD is configured with 3 network interfaces, one for each RPG and a separate one for network output to the SD's. Each SD is equipped with 2 network interfaces, one for each RPG. Thus each SD receives 2 versions of each product, one from each SD. If one RPG or network fails, the SD's continue to receive products from the other one.

D.4 Data Flow for Ribbon Display Generation

The primary function of the TDWR processing is to output ribbon display messages to controllers (accompanied by audio warning tone).

The standard IRIS output and display features are also used to construct images for the Situation Displays (SD's), that is, simplified displays showing runway warning areas (arenas), alert centroids and shear lines with forecast positions. Since the SD's are normal IRIS displays, other IRIS products can also be displayed.

Ribbon displays are generated when the local IRIS (the IRIS where the ribbon display is connected) receives a special TDWR product.
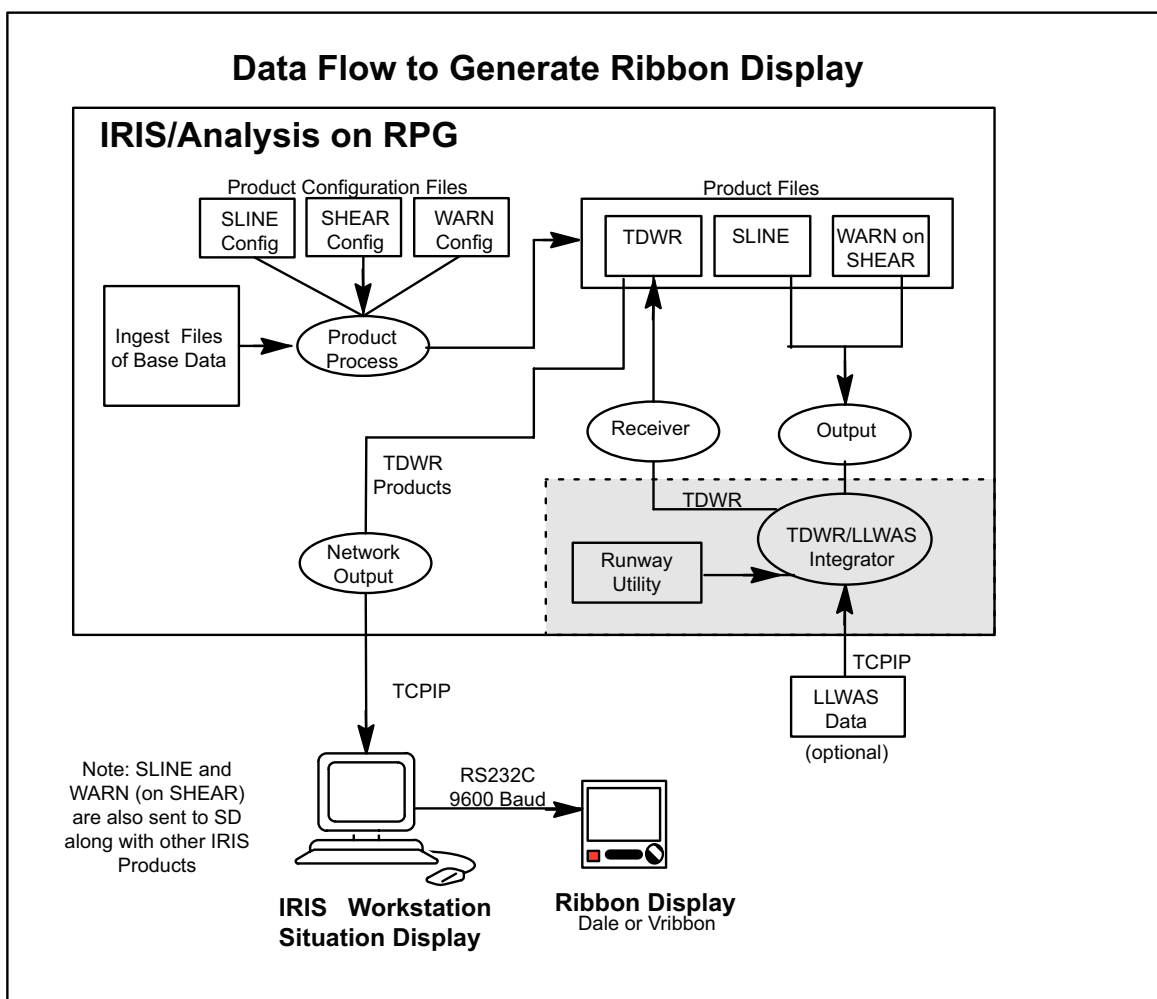
Figure 118   Data Flow for Generating TDWR Ribbon Display Products

Normal IRIS product generation makes **SLINE** and **WARN** products for radial shear (microburst warnings). These products have all possible protected areas turned on, even for inactive runways.

The results are sent to a separate TDWR/LLWAS Integrator. The Integrator takes data from LLWAS (if available) and configuration information on the current active runways from the runway utility. The output back to IRIS is a TDWR product that is filtered so that only the information required by the ribbon display is present. The integration algorithm is the FAA/MIT Lincoln Lab algorithm that was developed for the U.S. FAA TDWR.

Usually both IRIS and the Integrator run on the RPG, although it is possible to run the TDWR/LLWAS Integrator on an SD. This would, however, lead to configuration chaos since there are likely to be many SD's.

The RPG then outputs the TDWR product to the SD's and when it is received, it is automatically formatted into a ribbon display message and sent out the SD serial line.

The process that outputs the products to the RBDT's has an additional configurable filter: The controller can specify which corridors are displayed on an SD. Use the **ribsetup** to configure this.

D.5 IRIS/TDWR Configuration Summary

You must install and configure the TDWR components by performing the following procedures.

1. Configure IRIS for TDWR operation.

   a. Configure the runway arenas (protected areas).
      See D.6.1 Configuring Protected Areas (Arenas) for TDWR (page 351).

   b. Configure the scan tasks and the **SLINE** and **SHEAR** products.
      See D.6.2 Configuring IRIS Tasks and Products for TDWR (page 353).

   c. Configure the outputs to the Integrator and the SD's.
      See D.6.3 Configuring IRIS Product Outputs for TDWR (page 353).

2. Install and configure the Ribbon Displays. See:
   - D.7 Ribbon Displays (page 355)
   - D.7.5 Configuring Ribbon Display Software (page 360)

3. Configure the TDWR/LLWAS Integrator (`runways.conf`).
   See D.8 TDWR/LLWAS Integrator and Runways Utility (page 366).

4. Use the **runways** utility to set the runway direction.
   See D.8.3 Runways Utility (page 373).

D.6 IRIS Preparation

Preparing the IRIS system involves setting-up the protected areas, scan strategy, product generation, and output for TDWR operation.

D.6.1 Configuring Protected Areas (Arenas) for TDWR

Protected areas are configured in the IRIS **Setup > Product** utility.
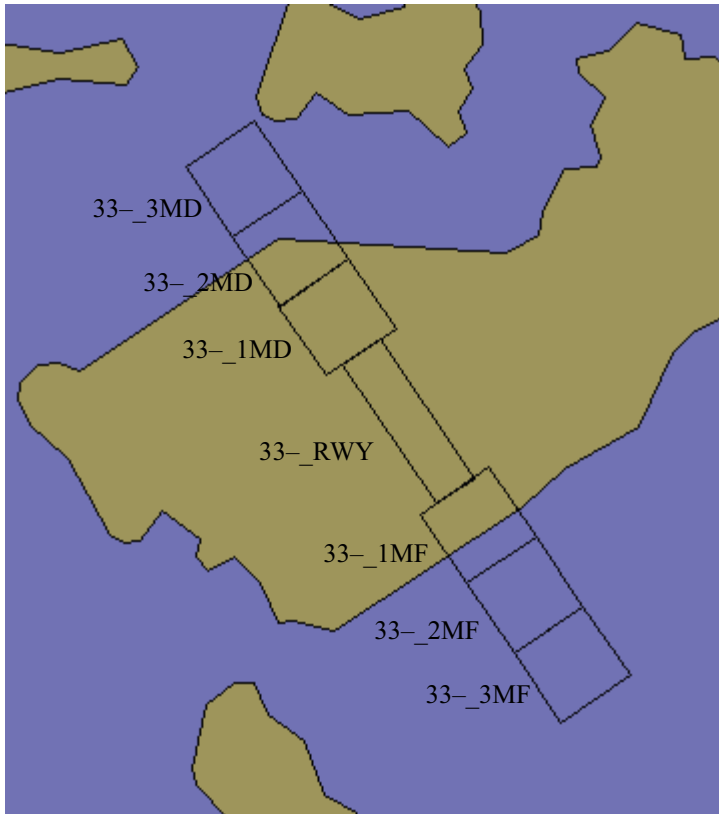
Figure 119   Configuring Protected Areas for TDWR

Arena Naming Convention

The naming convention used to designate runway corridors and protected areas (also called "arenas") is as follows:

- The first 2 digits are the runway name, also referred to as the "corridor" name. In the example, this is "33".
  Corresponding names for the reverse direction (runway 15) are made automatically. Vaisala recommends that you use the convention of naming the "north bound" direction for a corridor.
- The next position is used if there is a left and right runway. For example, use `33R` and `33L` to indicate left and right runways.
  If there is only a single runway, use a dash as in the example, that is, `33–`.
  Note, if the left and right runways are within approximately 500 m (1640 ft 5 in) of each other, it is easier to designate a single corridor for alert purposes rather than have 2 separate corridors that are almost completely overlapped.
- The next field is an underscore character which serves as a separator.
- The last 3 digits are used to specify the "arena". In the example, "`1MF`" specifies a protected area that extends from the end of the runway to 1 mile from the end of the runway. It is recommended that you use the text `RWY` to specify the runway itself.

Once you have decided the runway and arena names, construct the geometry of the protected areas, that is, the size and orientation information that is required in the IRIS **Setup > Product** utility. See *IRIS and RDA Utilities Guide*.

> It is tricky to get the areas correct. To help you Vaisala provides an EXCEL spreadsheet that makes it easy to configure runway protected areas. Contact your Vaisala representative for a copy.

After you have created your protected areas, test them with the **Overlay** utility. This requires that you have already configured the overlay for your system. See **Overlay** in *IRIS and RDA Utilities Guide*.

To see your protected areas on top of your overlay select **Overlay > Options > Protected**.

The display shows you the protected areas. Zoom in as needed to see them. The areas are not labeled.

Verify that all the areas are present in the correct locations and that all of the edges are properly aligned. If not, quit **Overlay** and then go back to IRIS setup/product to edit and save the protected area information. Retest with **Overlay** to see if you got it right.

D.6.2 Configuring IRIS Tasks and Products for TDWR

Configuring the scan tasks and the **SLINE** and **SHEAR** products required for TDWR operation depends on a the specific requirements of the user and the constraints of the site.

For example, some customers may require that a complete volume scan information be collected during TDWR operation (to provide general weather monitoring) while others may want to have the scan dedicated to low-level wind shear detection.

Most installations configure the system to switch automatically from a weather monitoring mode to a hazard monitoring mode when there is severe weather near the terminal.

Use the example configuration as a starting point for operation. You must adjust this example to meet your site requirements. This is typically done when your system is installed.

D.6.3 Configuring IRIS Product Outputs for TDWR

The IRIS product output configuration for the TDWR consists of:

- Output of TDWR products to the SD's (IRIS/Display Workstations with hardware and/or virtual ribbon displays).
- Output of **SHEAR** and **SLINE** products to the TDWR/LLWAS Integrator.

Both output configurations are made on the system where the TDWR/IRIS Integrator is running. This is typically an IRIS/Analysis Workstation functioning as an RPG (radar product generator). In the examples below, this node is called `cyclone`.

▶ 1. Configure output to the SD's:

    a. Setup the network output devices for the SD's in IRIS setup/output.
    The following example is for a network output to a workstation called **SD1**. The standard IRIS target directory and TCPIP notification port are used in the example.

    b. Configure the **IRIS Product Output Menu** to make automatic output of TDWR products to the SD's.
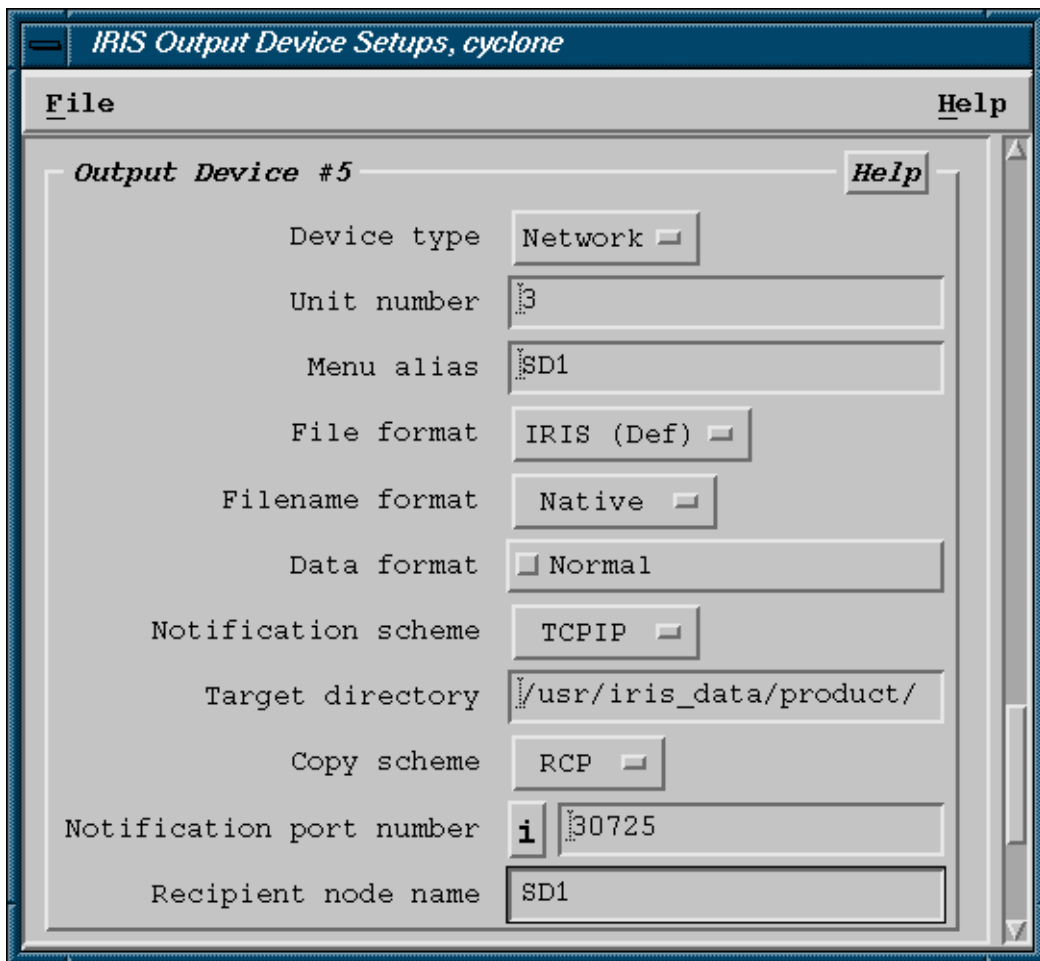
See 7.4 Sending a Product to a Device (page 274).

2. Configure output to the TDWR/LLWAS Integrator:

a. Use IRIS setup/output to configure the output to the Integrator.
   See the following example. Use the **Short** file name format.
   Since the Integrator runs on the same machine, use **Copy** rather than **RCP** (remote copy).
   The notification port must match that configured for the Integrator in the *runways.conf* file (See D.8.1 Configuring Integrator and Runways Utility - runways.conf (page 367)). The recommended port is 30726 as shown in the example.
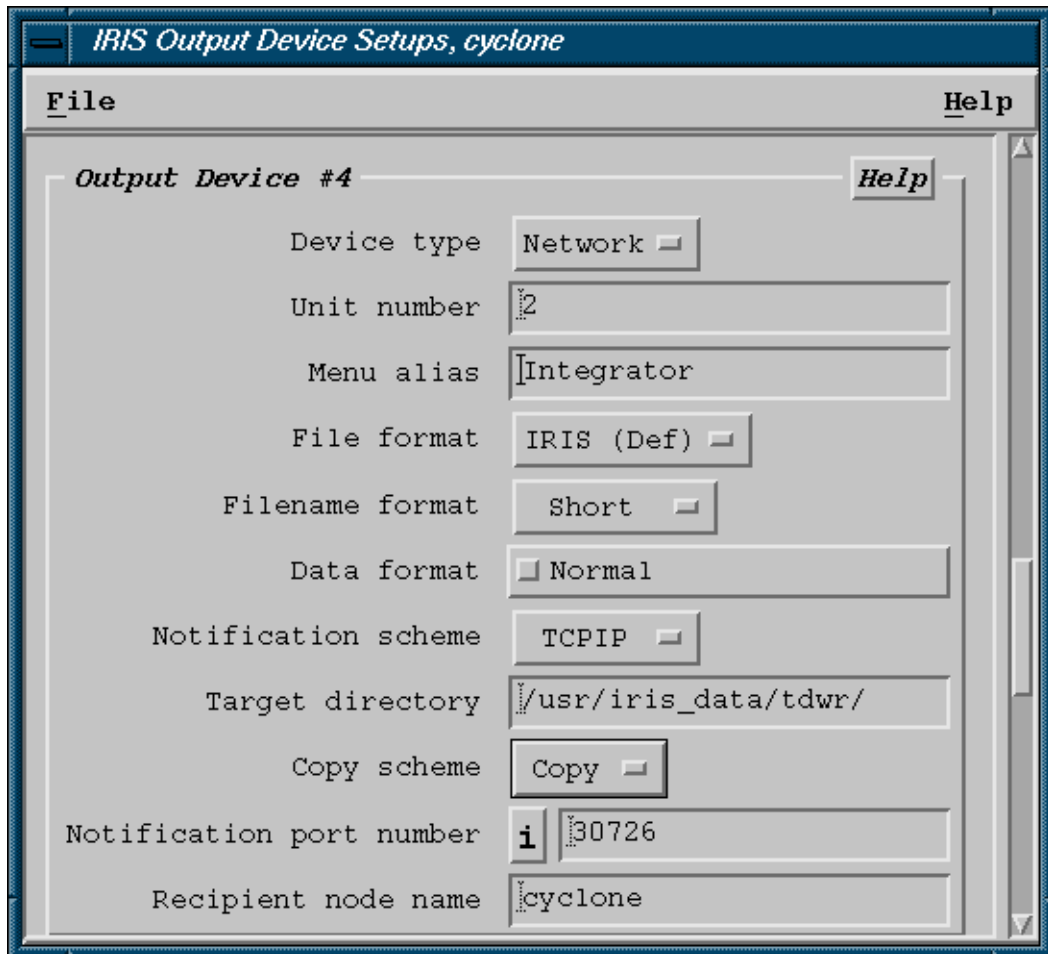
b. Configure the **IRIS Product Output Menu** to make automatic output of **SHEAR** and **SLINE** products to the Integrator.
   See 7.4 Sending a Product to a Device (page 274).

D.7 Ribbon Displays

A busy ATC controller does not have time to look at color situation displays and interpret them.

Ribbon displays are alphanumeric displays (connected to the Situation Displays) that provide ATC personnel with concise information on the status of various runway corridors with regard to wind shear and microburst events. The display shows simple text messages of the form:

```
32LA MBA 3MF 40K- 180 14
32LD 190 15
33LA WSA RWY 25K+ 190 17
33LD WSA RWY 25K+ 160 14
```

- The first line informs the controller of a microburst alert for runway 32 left at 4.8 km (3 miles) final approach. A 40 knot loss of airspeed can be expected. At the end of the line is the measured surface wind (direction and speed in knots) at the threshold end of the runway or from some other point on the field (depending on the availability of wind information from an LLWAS or other source).
- The second line shows the normal display with no warning message. Only the wind is displayed.
- The third and fourth lines show a wind shear alert (shear line) for runway 33L (approach and departure) with an expected 25 knot gain. The threshold winds are displayed at the end of each line.

In addition to the text messages an audio alert tone is sounded when an alert is issued. The controller acknowledges the alert tone by pushing a button on the front of the ribbon display.

D.7.1 Supported Ribbon Displays

IRIS Supports the following ribbon displays:

- Dale Electronics, Inc. Model ADP-250M045 (-1 for 220 V) Plasma Display Monitor/Alarm System. This is used by the US FAA. Because of its brightness it is well-suited for tower cab environments which have high background light.
- IRIS Virtual Ribbon Displays. This is an X-Window display which can be put on a workstation or PC running X-Windows.

The IRIS Workstation must be licensed for ribbon display support.

D.7.2 Dale Ribbon Display Hardware Installation

Up to 8 Dale ribbon displays can be "daisy-chained" to an IRIS workstation as shown in the following figure.

The **M** and **F** notation refers to the cable end (DB25M or DB25F).
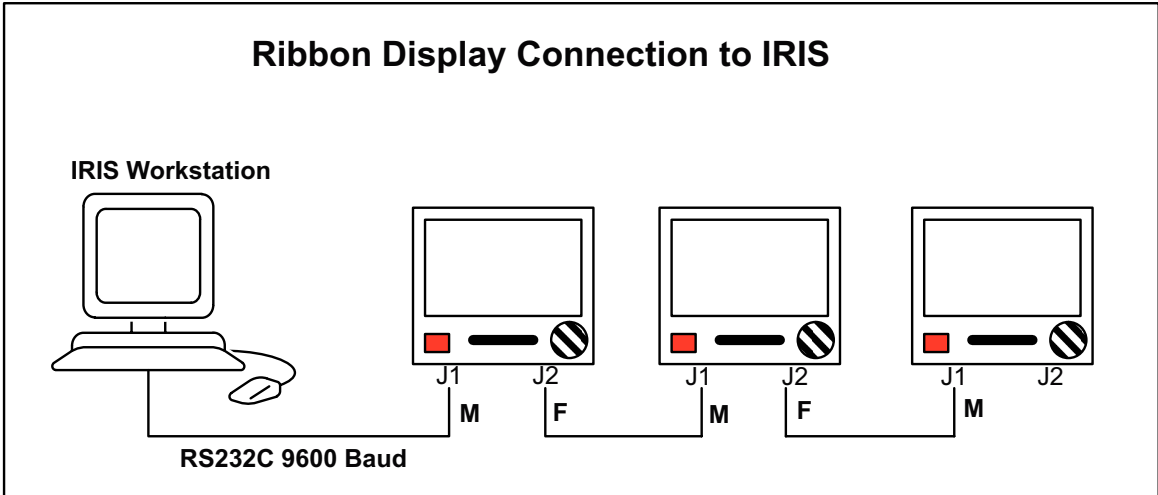
Figure 120   Dale Ribbon Display Installation

Table 26   Ribbon Display Pin Assignments

J2	J1	Signal
1	1	GND (shield)
3	2	Receive Data
2	3	Transmit Data
7	7	GND

D.7.3 Dale Ribbon Display General Use
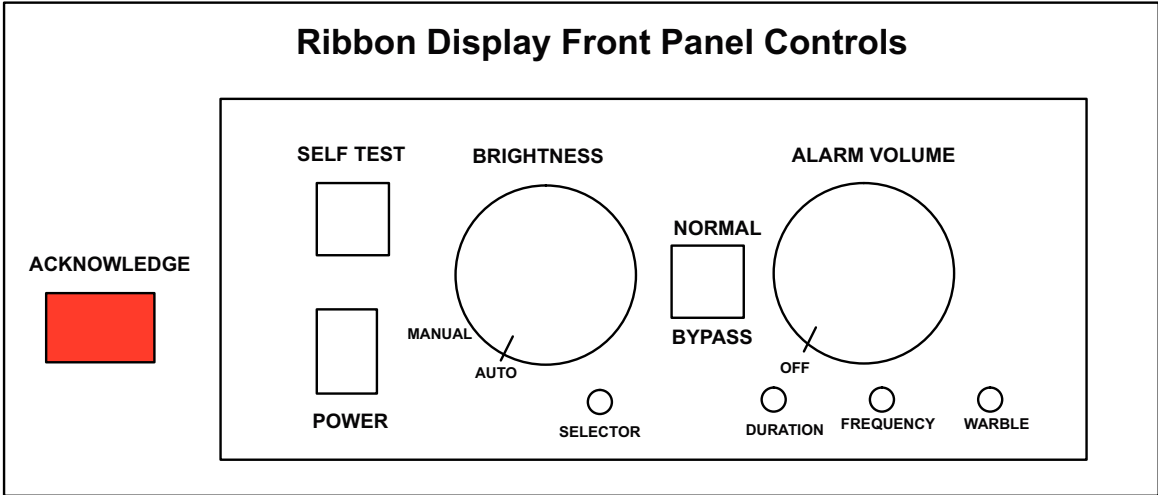
Figure 121   Dale Ribbon Display Controls

The hinged front panel drops down to access the local controls. The controls area is illuminated when the door is open.

Self Test

Hold in the **SELF TEST** push button to write the character set to the display. The monitor's ID number, model number, and software version are shown on the last line of the self test screen.

If required, you can use the selector switch can be adjusted to change the ID number. The audio alarm is activated and you can adjust the volume and warble.

If data is sent to the monitor while **SELF TEST** is depressed, the monitor sets D6 and D7 (of the status byte) high when the button is released to indicate that data sent to the unit was not captured and that the screen should be regenerated.

Brightness

Use the **BRIGHTNESS** control to adjust the display brightness.

To increase brightness rotate the control clockwise.

When the control is rotated to its extreme counter-clockwise position (past a switch detent) the brightness is controlled automatically according to the amount of ambient light in the room.

Normal/Bypass

Put the switch in the **NORMAL** position to enable communication between the monitor, host, and other peripherals.

Put the switch in the **BYPASS** position to route the RS-232D signals are routed around the monitor.

This hardware switch isolates the monitor's I/0 circuitry when in the **BYPASS** position.

Selector

You can set the monitor's identification number can be set from 0 ... 7 through the recessed switch.

To set the identification number, push and hold the SELF TEST button, and adjust the **SELECTOR** control until the desired number is displayed on the screen. Note: There are three positions that display 0. All work as ID #0.

Alarm Controls and Adjustments

The audible alarm is 2 tones that are alternately turn on and off to create a "warble" effect.

You can adjust the parameters of the warble using the **DURATION**, **FREQUENCY**, and **WARBLE** adjustments.

Volume

Use the **VOLUME** control to adjust the audible alarm loudness between 60 ... 85 dBa. Turn the control fully counter-clockwise to disable the alarm speaker output. Loudness is increased with clockwise rotation and decreased with counter clockwise rotation.

Duration

Varies the on/off time (duty cycle) of the two tones.

Frequency

Raises or lowers the frequency of the 2 warble tones.

Warble

Varies the rate at which the 2 tones are alternated.

Acknowledge
> When the host computer sends an alarm, the monitor activates the audible alarm and illuminates the **ACKNOWLEDGE** push-switch.
> Push the **ACKNOWLEDGE** switch to acknowledge and silence the alarm.

D.7.4 IRIS Virtual Ribbon Display

For systems that do not have the Dale model ribbon display, the IRIS Virtual Ribbon Display (`vribbon`) is provided.

It is also convenient for "eaves dropping" on Dale displays to verify that they are properly configured, that is, you can look at a virtual version of the Dale display.

Observers, Operators, and Controllers can access the `vribbon`. The vribbon display is run locally on an SD by opening a terminal and typing:

```
$ vribbon
```
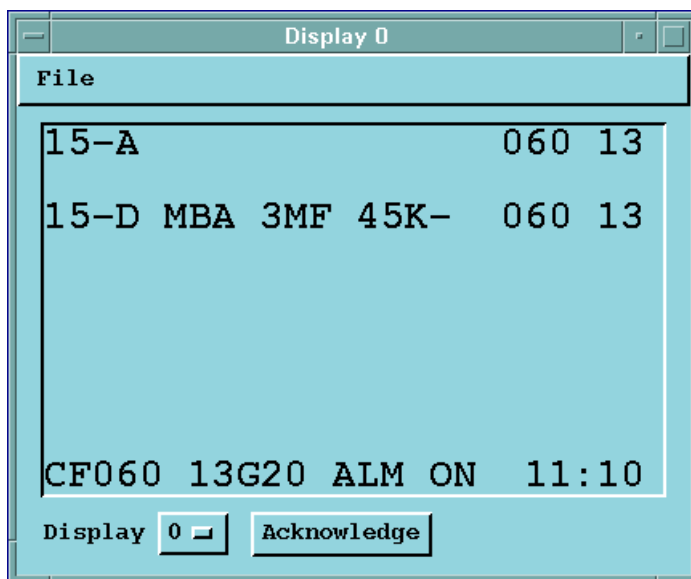

Figure 122   `vribbon` Display

The display has 10 lines. The top nine lines are used to display warnings for various approach and departure corridors. In the example, runway 15 departure has a microburst alert with a 45 knot loss. The threshold winds are reported as 60° (`060`) at 13 knots. The bottom line shows the center field wind data and the status of the alarm signal.

The `vribbon` can show any of the 8 physical ribbon displays. This is done by selecting the display number (0 … 7). The `vribbon` display also has a tone alarm that can be silenced (not disabled) by selecting **Acknowledge**.

D.7.5 Configuring Ribbon Display Software

ℹ️ To run the IRIS software for ribbon display support, your system must be licensed for this feature and you must have operator privilege. Check with your system manager.

Overview

The ribbon display setup utility provides the following functions:

- Runway approach and departure corridors are assigned to each line of the display.
- Up to 8 displays can be configured for each IRIS workstation and the lines for each display can be set independently.
- Up to 4 configurations can be saved. This is used for installations where different active runway assignments may require different display layouts.

Configure Ribbon Displays Using the ribsetup Utility

After you have defined your protected areas, the **ribsetup** utility is used to define the appearance of the ribbon display, that is, what runway information is displayed on each line of each ribbon display.

To start the **ribsetup** utility you must be an operator. Open a terminal window and type:

```
$ ribsetup &
```

The **ribsetup** utility appears as shown below:
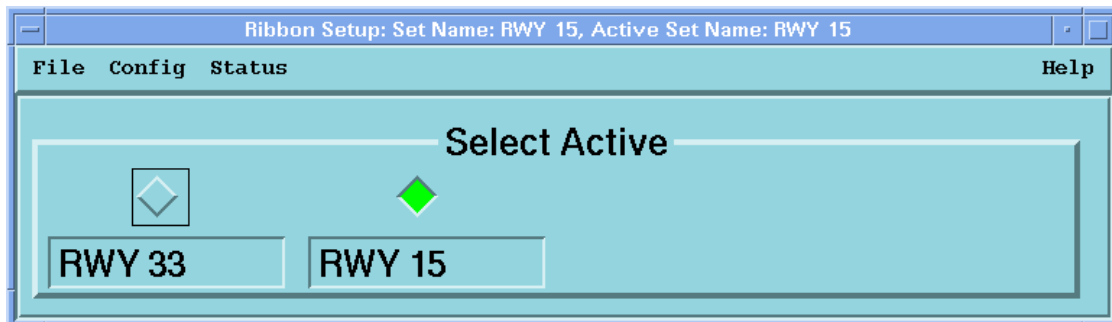
When the menu first comes up it is in "select" mode. You can select among up to four different configurations that have been pre-defined. In the previous example, 2 configurations are defined-one for runway 33 and the other for runway 15. The other 2 configurations are unused.

The configurations should be made to correspond to the actual runways that are typically used together in operation. When the terminal area is reconfigured in response to a new wind direction, this menu is used to select the appropriate ribbon display for the new configuration.

The features of the menu are described in detail below.

File

This provides the functions

- **Load**

Used to restore the menu from the currently-saved file configuration if you want to "undo" changes.

- **Save**

Saves your changes.

- **Print**

To print the menu for documentation purposes.

- **Exit**

To exit the menu. Note that ribbon displays remain active when you exit.

Status

Select **Status** to see a submenu of the communication status of each of up to 8 ribbon displays that can be attached to your IRIS workstation.

The top row of the menu shows the ribbon display numbers for 8 displays that could be connected to your IRIS workstation. Note that your system may have fewer than 8 displays. The second row shows the names that have been configured for the displays. The last four names are set to "none" to indicate that these are unused.

The status of each ribbon display is indicated as OK or Dead. A Dead display may have been turned-off or have a cabling problem. Also, the display number set on the display itself may be incorrect (see D.7.3 Dale Ribbon Display General Use (page 357)). The update time shows the time in seconds since the last communication between IRIS and the display. Input and output count are shown to indicate the number of bytes input from the display and output to the display. This is sometimes useful for debugging communications problems.

Config

The submenus under config provide the ability to customize the appearance of each ribbon display.

Config/Config … (serial output port)

This selection determines which serial line output port is used for ribbon displays. Only a single port is needed since the displays are daisy-chained together. Type the device name in the form:

/dev/ttyXXX

The **XXX** depends on your system, for example:

- Linux `ttyS0` or `ttyS1`
- HP `tty0p0` or `tty0p1`

Check with your IRIS System Manager if you are unsure.

Number of Active Sets (1, 2, 3, 4)

This determines how many different active sets you want. In the example, there are two active sets. You can name active sets by typing in the name that you want and selecting **File > Save**.

Display Setup Menu

This is turns on the setup menu. The setup menu is used to configure the ribbon display appearance for each active set.

The display configuration for an active set can be viewed by pushing the button under the active set name. Note that this does not change the active set on the ribbon displays, that is, it only effects what is displayed and edited in the setup menu. The button turns yellow to indicate which set is being viewed/edited. In the example, the setup menu is set to view `RWY 15` (yellow button).

The top of the setup menu shows the display number and is color coded red or green to indicate whether display status is `OK` or `Dead`. The display numbers correspond to the "Selector" position (0 to 7) that is configured on each of the ribbon displays. It is recommended that if, for example, you have 4 displays, that the selector switches on the displays be set to 0, 1, 2, and 3.

There is also a button next to the Display Number to select whether a display is enabled or not. If a display is installed, the button should be clicked-in. In the example, the first four displays (0 to 3) are enabled, but displays 1, 2 and 3 are dead (read color) indicating that they are either turned-off or there is a communication problem.

Belowh each display number is a display name (for example, `Tower1`). Set these to indicate the location or function of the particular ribbon display. It is recommended that you use the text `None` or `Unused` to indicate the case when a display is not installed. To change the name, type the text.
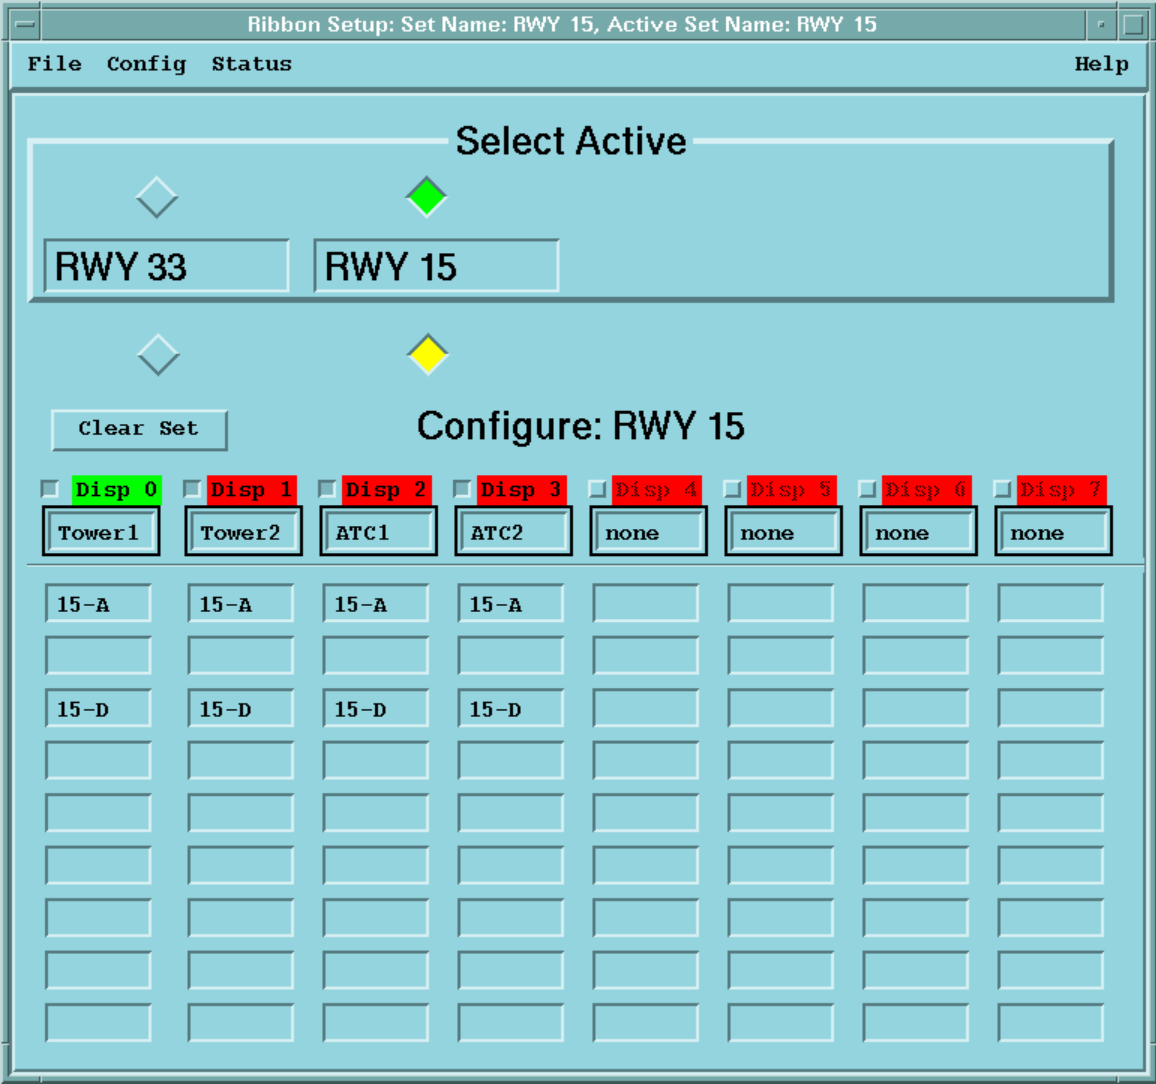
Figure 123   Ribbon Display Setup Menu

The bottom of the setup menu has 9 entries for each display, corresponding to the 9 lines that are available for warnings on a ribbon display. To configure what appears on each line, first highlight the line and then right-click the mouse to get a pop-up list of all possible corridors as shown below.

The corridor names are taken from the IRIS setup/product warning areas that were created (see the beginning of this section). In addition, there is a blank line (to configure blank lines on the display), and there are other arbitrary protected areas that may not be associated with a particular corridor. For example, a single large protected area could be configured to display alerts anywhere within 8 km (5 miles) of the terminal.

Select the corridor from the submenu for each line of the display. Repeat the procedure for all of the displays.

In the example, each of the four displays that are enabled is configured identically. In general this is good practice since it makes the system easier to maintain. If your operation requires it, different displays can have different configurations.

To assist with major changes, use the **Clear All** button. If you make a mistake, use **File > Load** to restore the saved values.

After you have configured a set, use **File > Save** to save your results. You may also want to do a **File > Print** to document your settings. Then you can click on a different set (yellow button under the set name) and configure the displays for this. Continue until you have configured all of the active sets.

D.7.6 Ribbon Display Testing with `tdwr_sim`

> ✋ **WARNING!** Do not run the `tdwr_sim` test utility during normal system operation as it may interfere with the signaling of actual warnings.

To test the ribbon display and for training purposes, Vaisala supplies the `tdwr_sim` utility. The utility outputs a simulated TDWR product with alerts for selected corridors to verify that the ribbon displays are functioning. Note that this does not verify that warnings are properly generated by the IRIS system.

▶ 1. Start a terminal and type:

```
$ tdwr_sim
```

2. When prompted to `Enter Corridor Count`, enter the number of corridors that you have on your system.
   In the example that has been used in this appendix there would be four (**33A**, **33D**, **15A**, **15D**). The maximum count supported by the simulator is 16.

3. Press **ENTER**.
   The corridor names are displayed line-by-line. You can type arbitrary text to generate a test message.
   A sample session is shown below:

```
Enter Corridor Count [4]: 4
Enter Alert type for Corridor 0 (15-A)[ ]:
Enter Alert type for Corridor 1 (15-D)[ ]: MBA
Enter Alert type for Corridor 2 (33-A)[ ]:
Enter Alert type for Corridor 3 (33-D)[ ]:
Sending product to IRIS.
Enter Corridor Count [4]:
```

`tdwr_sim` sends a simulated TDWR product to IRIS, which is transmitted to each ribbon display that is configured to show alerts for the selected corridor (in this case 15-D).
This causes the ribbon display to have a warning message as shown below.
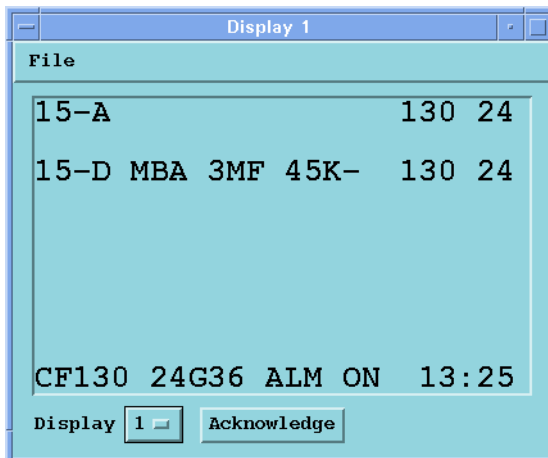
4. To clear an alert, press the SPACE BAR to replace the warning with a blank line.
   If no new message is sent to the ribbon displays for 150 seconds, the text `NO DATA AVAILABLE` is shown on all ribbon displays.

5. To stop `tdwr_sim`, type **CTRL+C**.

D.7.7 Ribbon Display Alert Messages

33–A MBA 3MF 40K– 180 14

Approach/Departure Corridor
Type of Alert (Microburst or Wind Shear)
Location (e,g, 3 miles final approach)
Expected airspeed gain/loss
Wind direction at threshold or other specified location

Figure 124   TDWR Ribbon Display Alert Messages

D.8 TDWR/LLWAS Integrator and Runways Utility

The TDWR/LLWAS Integrator is a separate software process (that is not part of the main IRIS software) which combines wind shear alerts from IRIS and an optional LLWAS system to generate TDWR products. The Integrator functions as a filter on the alerts, that is, the integrator:

- Only passes alerts for runways that are active as specified by the runways utility.
- If an arena has 2 or more alerts, then the more important alert is passed, for example, a microburst alert is more important than a wind shear alert.
- Only passes the alert for the arena of first encounter in a corridor. For example, if **33L_3MF** and **33L_2MF** both have microbursts alerts, then the **33L_3MF** is alerted on the ribbon display since an aircraft would encounter this first.
- In the event of failure of a system, the failed system uses the alerts from the other system, for example, if TDWR fails, then only LLWAS data is used.
- For each arena, the alerts from TDWR and LLWAS are combined as follows: a strong alert (MBA) from either system is passed through, while a weak alert (WSA) requires confirmation from the other system. The thresholds for what is considered strong and weak are tuned for each site.

The TDWR products from the Integrator are sent back to IRIS for distribution to SD's to drive the ribbon displays. See D.4 Data Flow for Ribbon Display Generation (page 350).

The Integrator takes as input the following data:

- **SLINE** products from IRIS (gust front warning)
- **WARN** (on radial shear) products from IRIS (microburst warning)
- LLWAS input data (if available, for LLWAS warnings).
- Active runway configuration from the runways utility.

A new TDWR product is made and output to IRIS when:

- New input data are received.
- The active runway configuration is changed.
- A maximum time-out has expired (specified time of order 60 seconds).
- The product sitename is the same as the radar sitename.

All of the inputs are used when a TDWR product is made, unless an input is "stale", that is, it is older than a specified time (of order 120 seconds). For example, if LLWAS fails, the Integrator continues to use the old LLWAS data until the time limit for "stale" data is reached. If all of the data are "stale", no TDWR product is made at all. In this case, there is no output to IRIS and the ribbon displays eventually show NO DATA AVAILABLE.

IRIS/Integrator Communication
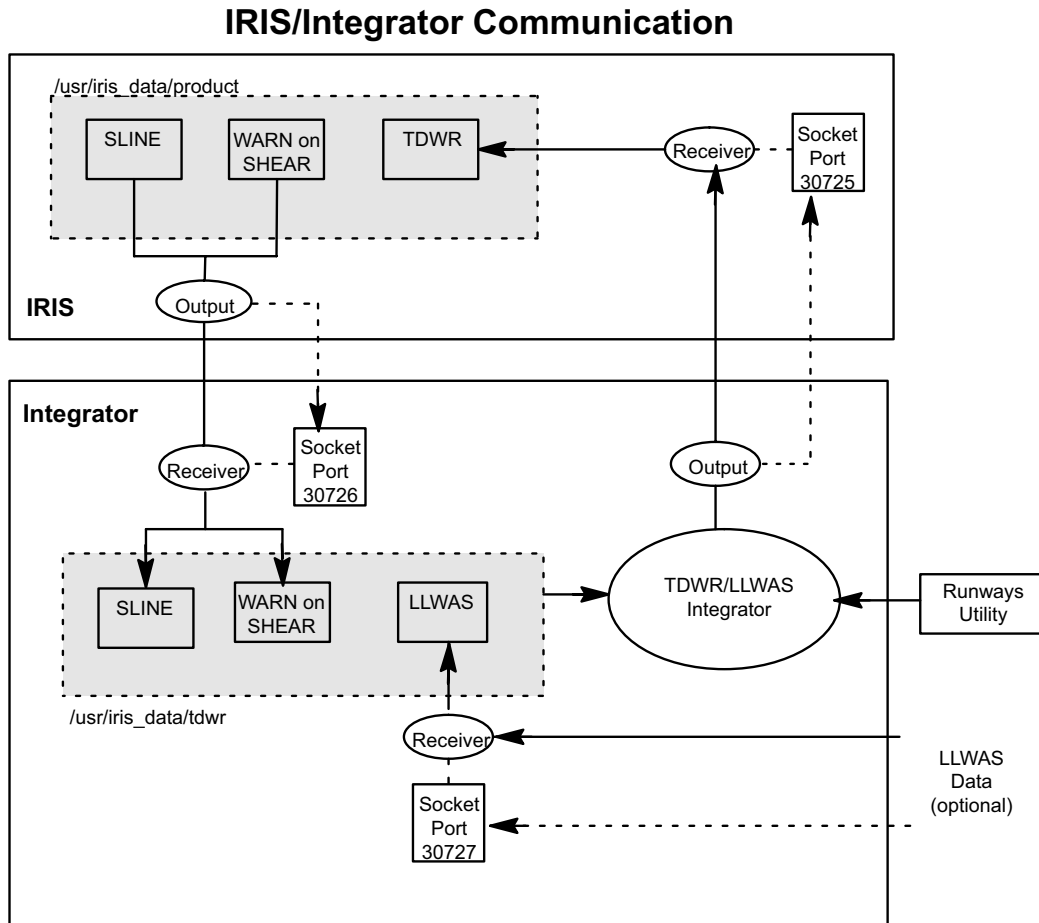
Figure 125   IRIS/Integrator Communication

D.8.1 Configuring Integrator and Runways Utility - runways.conf

The configuration of the integrator and the runways utility is stored in the *runways.conf* file in the *IRIS_CONFIG* directory which is assumed here to be the */usr/sigmet/config* directory.

The configuration steps are:

1. Create the *tdwr* data directory (default location */usr/iris_data/tdwr*)
2. Install the *runways.conf* file (default location */usr/sigmet/config/runways.conf*)
3. Edit the *runways.conf* file.

Create the directory */usr/iris_data/tdwr*

You must have root privilege for this step. The *tdwr* directory can exist anywhere on the workstation that runs the Integrator. However, we recommend that you put it in the same place that you put the *IRIS_PRODUCT* product directory.

For a standard installation this is in a directory called */usr/iris_data*. This may be different on your system. To see the location of the product directory check */usr/sigmet/config/profile*.

To make the *tdwr* directory, log in as root or become the super user and type:

```
# cd /usr/iris_data
# mkdir tdwr
```

Set the privileges for the directory to match the other directories (do an `ls -ll` to see how the other directories are set). For a standard installation type:

```
# chown operator tdwr
# chgrp users tdwr
# chmod 777
```

Install the runways.conf file

The *runways.conf* file is kept in the */$IRIS_ROOT/config* directory (usually */usr/sigmet/config*).

Check if the file exists in this directory. Log in as operator and type:

```
$ cd /usr/sigmet/config
$ ls
```

Check that the *runways.conf* file is in the list. If it is not, copy the template as follows:

```
$ cp /usr/sigmet/config_template/init/runways.conf /usr/sigmet/config
```

Edit the runways.conf file

Edit the *runways.conf* file (use `vi`, `emacs`, or your ASCII text editor). The file is shown below:

```
# File: ${IRIS_CONFIG}runways.conf
# Contains configuration information for the TDWR/LLWAS integration
# algorithm  and the runways utility.
# Number of runways to define: COUNT 3
# For each runways: Primary direction, Secondary direction, default direction.
# The default direction is the direction to use at powerup until changed
# manually  by runways.

RUNWAY 1 33- 15- LLWAS

RUNWAY 2 36R 18L LLWAS

RUNWAY 3 36L 18R LLWAS

# Uncomment this if you want all corridors reported out of integration.
#DONT_FILTER_RUNWAYS
# Turn logging on, choices are FILE or TERM or BOTH or NONE LOG BOTH
# Define input related info:
# Socket port; Directory; Timeout to consider input data stale

IRIS_IN 30726

IRIS_IN_DIR /usr/iris_data/tdwr/

IRIS_IN_TIME 120

# Define output related info:
# Socket host and port; Directory;
# Timeout to retransmit output even with no changes

IRIS_OUT cyclone 30725

IRIS_OUT_DIR /usr/iris_data/product/

IRIS_OUT_TIME 60

# Define LLWAS input related info:
#LLWAS_IN 30727
# Since the LLWAS input does not contain sitename or lat/lon. These
# need to be specified to handle the case where IRIS input is lost.
# For systems with no LLWAS input, it is recommended that you comment
# this whole section off.

# Specify radar lat and lon in signed degrees. Negative is south latitude, or
# West longitude.

RADAR_LATLON 43.00 -71.00

# If you want to take the latlon from the IRIS input, if  available, then
# comment out this line.

ALWAYS_SET_LATLON
```

```
  # Specify desired radar sitename. Max 16 chars, requires quotes if you
  # have imbedded whitespace.

  RADAR_SITENAME "THIS IS A TEST"

  # If you want to take the sitename from the IRIS input, if available, then
  # comment out this line.

  ALWAYS_SET_SITENAME

  # List of users who are allowed to operate the runways utility.

  OPERATORS operator

  # End of file
```

The lines in the file are described in detail below:

- ```
  COUNT 3
  ```

  This is the number of physical runways, that is, strips of concrete.

- ```
  RUNWAY 2 36R 18L LLWAS
  ```

  For each runway, an arbitrary number is assigned (for example, 1, 2 and 3). These
  numbers should be sequential starting from 1.
  The names for each runway are then specified in the next two fields (for example, `36L`
  and `18R`).
  The last field is the default active runway when the software is first started. This can be
  either of the runway names or, if available, LLWAS. In the LLWAS case the runway
  direction specified in the LLWAS data is used when the Integrator is first started. If
  there are no LLWAS data, then the runway direction in the first field is used on start-up.

- ```
  # DONT_FILTER_RUNWAYS
  ```

  This line should generally be left commented (the # sign).
  If it is un-commented (# sign removed), the integrator passes on alerts for all arenas
  even if they are not associated with active runways.
  Filtering can still be performed by the ribbon displays, that is, `ribsetup` can be
  configured to exclude inactive runways. See D.7.5 Configuring Ribbon Display Software
  (page 360).

- ```
  # Turn logging on, choices are FILE or TERM or BOTH or NONE
  LOG BOTH
  ```

  - If you are starting the integrator from a terminal, set this to **BOTH**.
  - If you are starting the integrator automatically, set this to **FILE**. The log file is stored in */usr/iris_data/log*.

- ```
  IRIS_IN 30726
  IRIS_IN_DIR /usr/iris_data/tdwr/
  ```

These entries configure the characteristics of the input from IRIS to the Integrator. `IRIS_IN` is the socket port that is used by IRIS to notify the Integrator when a new product is sent. This must match the IRIS setup/output network device that has been configured to send products from IRIS to the Integrator. The default port is 30726. The *IRIS_IN_DIR* is the directory where IRIS sends products to the Integrator. This must match the directory configured at the beginning of this section and the directory configured in IRIS setup/output for the TDWR integrator. See .

- ```
  IRIS_OUT cyclone 30725
  IRIS_OUT_DIR /usr/iris_data/product/
  ```

These entries configure the characteristics of the output from the Integrator to IRIS. `IRIS_OUT` is the nodename (cyclone) and socket port (30725) where IRIS will receive the data. Usually this is the local workstation where the Integrator is running. The socket port should match the `IRIS_NETRCV` environment variable for the network receiver. This is set in the */usr/sigmet/config/profile*.
The *IRIS_OUT_DIR* is the directory where TDWR products are sent. This must match the directory configured for *IRIS_PRODUCTS*, typically */usr/iris_data/product*. This is set in the */usr/sigmet/config/profile*.

- ```
  IRIS_IN_TIME 120
  ```

This is the time in seconds, after which, an input is considered to be "stale". For example, if the LLWAS stops sending data, the Integrator continues to use the old LLWAS data for 120 seconds (in the example). After 120 seconds, the old LLWAS data is not used to make new TDWR products.

- ```
  IRIS_OUT_TIME 60
  ```

This is the maximum time between generation of TDWR products. A new TDWR product is made if one of the following happens:
  - A new input arrives.
  - The active runways configuration changes.
  - One of the inputs becomes stale (`IRIS_IN_TIME` times-out).
  - `IRIS_OUT_TIME` from the last TDWR product generation is exceeded.
  - The `IRIS_OUT_TIME` makes sure (in this example) that a TDWR product is created at least every 60 seconds provided there are valid inputs (data are not stale).

- ```
  LLWAS_IN 30727
  ```

This is the socket port that is used by LLWAS to notify the Integrator when a new LLWAS data are available. If you do not have and LLWAS, comment this line by adding # to the beginning.

- ```
  <LLWAS Site Information>
  ```

If you have an LLWAS system, you must specify the LAT/ LON of the radar and the sitename since these are not in the LLWAS header.
This information is needed if the radar fails and the alerts are based on LLWAS alone. If you do not have an LLWAS, comment all lines of this section.

- ```
  OPERATORS operator controller>
  ```

This controls access to the runways menu to the user ID's in the list. Note that anyone can view the runways menu, but only users in the list can change it.

D.8.2 Starting Integrator Utility

Manual Integrator Start-up

To start the TDWR/LLWAS Integrator you must be an IRIS Operator. In a terminal window on the workstation where the Integrator is run (typically an RPG IRIS/Analysis system) type:

```
$ tdwr_llwas_int
```

This starts the software and log events in the terminal window. Iconify the terminal so that you can read any messages in the log.

Automatic Integrator Startup

The recommended approach is to start the integrator automatically when the workstation boots-up. This is done by creating a startup file in a manner identical to the automatic **siris** startup.

D.8.3 Runways Utility

The **runways** utility is used by controllers to specify which runways are active. This information is used by the TDWR/LLWAS Integrator algorithm to filter alerts from inactive arenas. For example, each runway arena has two names, for example, 33L_3MD is the same physical arena as 15R_3MF. The runway direction specified in the **runways** utility determines which of these names is used to identify the arena. Also, if a runway is not used, the entire runway and all of its associated arenas may be set to inactive.

The **runways** utility must be run on the same IRIS node where the Integrator is run (typically this is the IRIS workstation assigned to generate products), however you can export the **runways** display to another workstation on the network. Thus the **runways** utility can be run from any networked workstation.

Whether you are on a local or remote workstation, you must have privilege to log on to the system. Note that most installations allow Observers to view the runways utility, but only Controllers (and perhaps Operators) can change the runway direction. The OPERATORS line in the *runways.conf* file specifies which users can control the **runways** utility. See D.8.1 Configuring Integrator and Runways Utility - runways.conf (page 367).

Starting the Runways Utility

The **runways** utility may be configured to start automatically on your system. This is configured by the system manager. See D.8.1 Configuring Integrator and Runways Utility - runways.conf (page 367).

In a command line terminal, type:

```
$ runways &
```

If you are on a different workstation you can use **rlogin <nodemame>** or **sigterm <nodename>** to open a terminal on the node where the integrator is run. Note that if you use the **rlogin** approach, before running runways you must type the following:

```
$ export DISPLAY=< yournodename >:0.0
```

This exports the runways display to your workstation. If the runways display does not appear on your workstation, then you may need to type the following on a terminal on your workstation (not the remote system):

```
$ xhost +
```

This authorizes windows from other workstations to appear on your display. Check with your system manager if you want this automatically enabled.

Stopping the Runways Utility

To exit the runways utility select **File > Exit** in the upper left. Note that the settings specified in the runways utility are used by the Integrator after you exit.

Setting Active/Inactive Runways

Each runway that is configured in *runways.conf* must have a line in the menu as shown at the beginning of this section. The button in the far left column activates/deactivates the runway. Inactive runways are indicated by the "X" symbol over the runway, as shown in the example for 33/15.

Selecting the Runway Direction

The runway direction is indicated by the green color of the runway name and the direction of the aircraft icon. The runway direction can be selected either manually by clicking the runway name, or automatically by LLWAS (if available). To enable the LLWAS automatic selection, click the button on the LLWAS column for each runway.

LLWAS Status
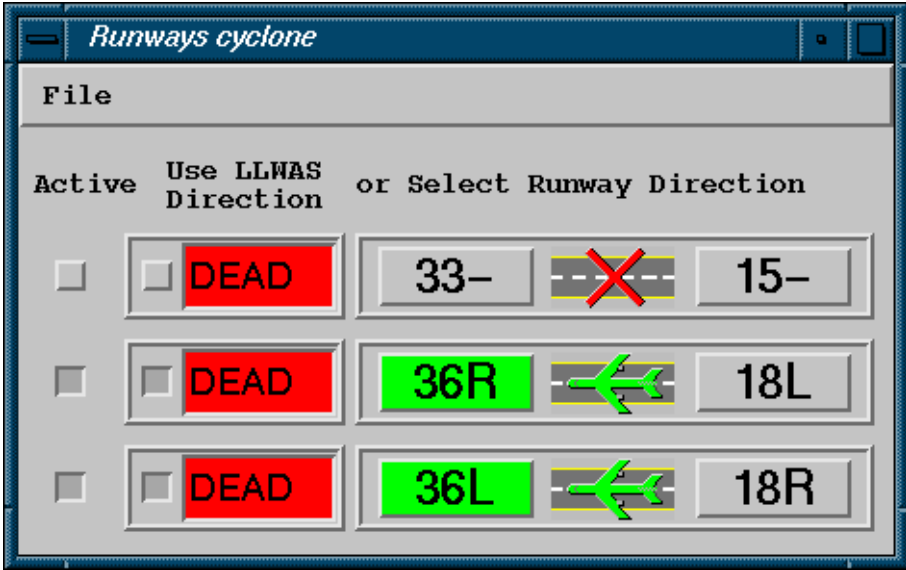
If LLWAS is available on your system, its communication status with the integrator is displayed in the LLWAS column. The image above shows the display DEAD in red background to indicate the LLWAS is not communicating to the Integrator.

Appendix E. Hydromet Raingage Correction

E.1 Hydromet Raingage Correction Overview

IRIS/Hydromet is an optional IRIS package for hydrometeorological applications.

Table 27   IRIS Features for Hydrometeorological Applications

Standard Features	Optional Features
Display of rainfall rate in mm/hr or inches/hour for a Z–R relationship selected in the **Product Configuration** menus for: • **PPI** • **RHI** • **CAPPI** • **MAX** • **XSECT** The units selection is made by the color scale selection.	**CATCH** product: • Provide precipitation average depth in subcatchment regions for selectable time periods. • Displayed as color-coded regions or as color-coded icons. Subcatchment regions are defined similarly to overlays. In the Quick Look Window, you can right-click the mouse on a region or icon to see a bar graph representation of the subcatchment amounts by hour.
RAIN1 Hourly rainfall accumulation for selectable Z–R relationship.	**GAGE** product input with raingage calibration. This allows the radar **RAIN1** hourly accumulations to be calibrated to the actual rain that is measured in gages. **GAGE** products may contain the **ZR** coefficient and exponent (by site) derived from disdrometer measurement or other source. See 3.2.3 Product Parameters (page 25).
RAINN Multiple-hour rainfall accumulation based on the hourly amounts.	View a 12-hour history.

More Information

‣ CATCH: Subcatchments Precipitation Accumulation (page 127)
‣ GAGE: Raingage Reports (page 172)
‣ RAIN-N: N-Hour Rain Accumulation (page 63)

E.2 Hydromet Data Flow for Raingage Correction

The following figure shows the data flow for the raingage correction.
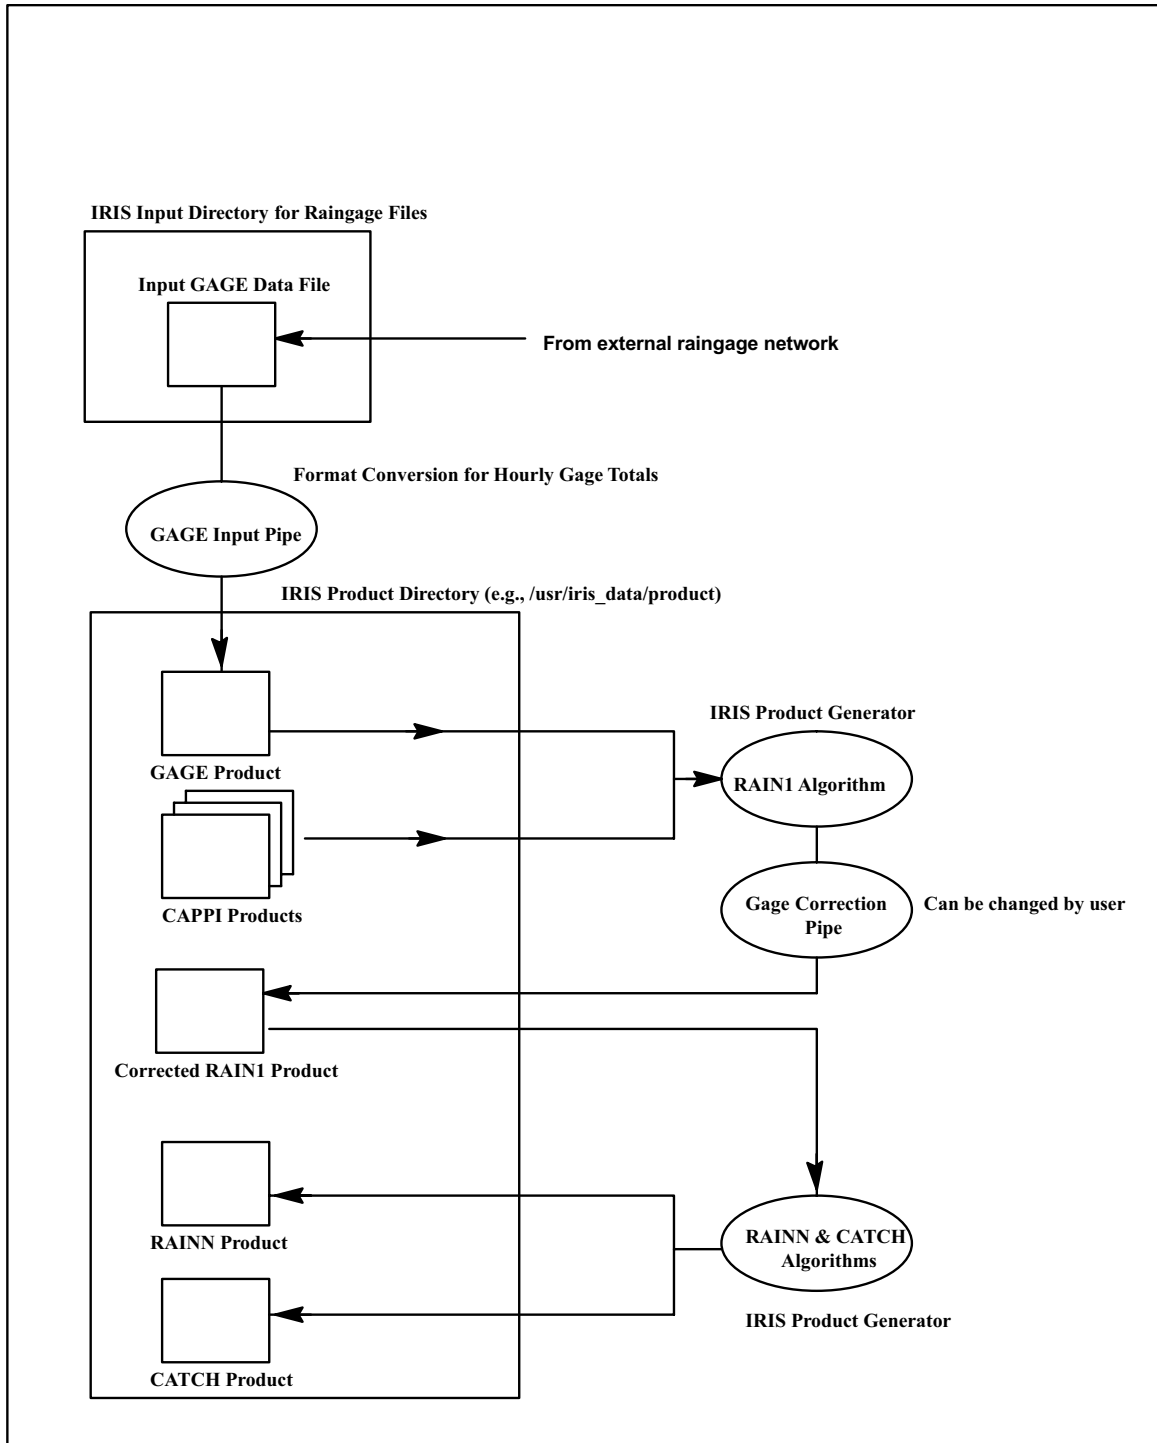
Figure 126   **RAIN1** Raingage Correction Data Flow

- Input from External Raingage Network
  Raingage data are supplied by the customer in the form of files that are transferred automatically by the customer over the network to a directory that is defined by the IRIS **Setup > Input**.
  Each file should contain the total accumulation for the previous hour for all gages. See E.3 Gage Data File Format (page 380).
- Input Pipe for Format Conversion to **GAGE** product
  The Input Pipe is a format convertor to reformat the customer data to the **GAGE** product format. The Input Pipe is specified in **Setup > Input**.
  It is a software program that is usually customized for the particular raingage network. You can create their own pipes or Vaisala can provide the pipe `AsciiToGage`, in which case the pipe source code is also provided so that customers can maintain the pipes in the future.
- **GAGE** Product
  This is a product in IRIS that contains the raingage totals for all gages in the network for the last hour. See E.3 Gage Data File Format (page 380).
- Input **CAPPI** products
  The IRIS **RAIN1** hourly rainfall accumulation product uses **CAPPI** products as input. The **CAPPI** products can be in **dBZ**, **dBT**, or the corresponding rainfall rate according to a default **Z–R** relationship such as $Z=200R^{1.6}$.
  Usually a "Filled" **pseudo CAPPI** is used so that at far ranges the lowest elevation angle is used to fill in the data.
- **RAIN1** Processing and Raingage Calibration Process
  Gage correction is activated when the gage correction button is turned on in the **RAIN1** product configuration menu. In this case, the IRIS product generator waits for all of the **CAPPI**s and the **GAGE** product for the previous hour and then runs the **RAIN1** algorithm.
  For information on the scheduling considerations for the **RAIN1** algorithm with **GAGE** correction, that is, how to handle the time delay of the arrival of the **GAGE** products, see E.6 RAIN1 Scheduling with GAGE Product (page 386).
  The RAIN1 algorithm proceeds normally as described in 3.12 RAIN1: Hourly Rain Accumulation (page 60).
  The extra step of the **GAGE** correction is done by a "Product Pipe" algorithm. Vaisala supplies an algorithm based on Brandes and Wilson (1982). You can supply your own product pipe correction algorithm in lieu of the one supplied. See *IRIS Programming Guide*.
- Corrected **RAIN1** Product
  The raingage calibration process outputs the final corrected **RAIN1** product for use by the **RAINN** and the **CATCH** products.

> You can make an uncorrected **RAIN1** product can be made by running the **RAIN1** product (under a different name) with the gage correction button disabled. For example, name the uncorrected product `HOURLY` and the corrected product `HOURLY- GAGE`.

- **CATCH** and **RAINN** products
  The **RAIN1** product serves as the basis for the **RAINN** and **CATCH** products.
  You can also run **WARN** products can be run on the **RAIN1**. Thus all subsequent
  products that rely on **RAIN1** for input benefit from the correction algorithm.

E.3 Gage Data File Format

The raingage data files should be sent automatically to a directory specified in IRIS **Setup > Input** for the input pipe.

The data can be sent by FTP, RCP, or NFS. The file should contain the hourly gage totals for all of the gages along with other supporting housekeeping information required by the **GAGE** product such as date and time and raingage location.

The input data file format is as follows:

- Each file is plain text ASCII, and the overall structure is a collection fields consisting of a keyword followed by one or more data values. This makes the format expandable, since new keywords can easily be added in the future, and each software reader can choose which set of keywords it needs to process.
- White space is used to separate the keywords and data values. Quotation marks may be used around a string value if there is a possibility that the string contains white space. Each line is terminated by a "newline" character.
- Each file represents data from any number of stations, but only one data time is reported in each file. That time is the ending time of the data, and is reported using the `TIME` and `SPAN` keywords in the first lines of the file. The fundamental time resolution is quantized to one-minute intervals.
- Each rain gage station report consists of a line containing at least the `CODE` and `LONLAT` keywords, plus either the `RRATE` or `RFALL` keyword. The `QUAL` and `REM` keywords are optional.

In general, the input pipe must be custom, written either by the customer or by Vaisala. This is the format used by the *AsciiToGage* input pipe. If you wish to use another format, you must write your own custom input pipe.

> This recommended input gage data format does not match the format of the **GAGE** product since the input pipe reformats the product for IRIS.
> The basic information must be in the file, or in some cases, in a separate configuration file for the input pipe. See *IRIS Programming Guide*.

The keywords are as follows:

`# any comment`
Any line beginning with at "#" is ignored, as well as all blank lines.

`TIME yyyymmddhhmm`
The ending time for all of the stations reported in this file. The string of digits gives the year, month, day, hour, and minute. This keyword appears once at the beginning of the file.

SPAN nn

The time span in decimal minutes for all of the data reported in this file. This keyword appears once at the beginning of the file. IRIS supports 60 minute spans.

CODE nnnnnn

Station code for a rain gage. The code is an arbitrary string 15 or less characters long. This field must be present for each reported gage.

LONLAT nn.nnnnn nn.nnnnn

Longitude and latitude of a rain gage. Longitude is a real number from -180 ... +180, with positive values representing East longitude.
Latitude is a real number from -90 ... +90, with positive values representing North latitude. Any number of decimal digits of precision can be used. This field must be present for each reported gage.

RRATE nn.nnn

The rain rate in millimeters/hour for a raingage. Either the **RRATE** field or the **RFALL** field must be present for each reported gage.

RFALL nn.nnn

The rain fall amount in millimeters for a raingage. Either the **RRATE** field or the **RFALL** field must be present for each reported gage.

Z/R ccc.c e.ee

Optional field that is available from disdrometer type raingages capable of calculating a **Z/R** relationship. The numbers are the constant and exponent in the equation $Z = CR^E$. When there is no rain, or the disdrometer is unable to make a measurement, the values of 0.0 are sent for both numbers. This is interpreted as a special flag to indicate that the default values should be used.

QUAL nn

The data quality (0:Useless, 10:Best) for a rain gage report. This field is optional. If it is omitted the data quality is assumed to be "10" (Best).

REM xxxxxxxxxxx

Allows a comment to be inserted in the file. An arbitrary text field follows. Use quotes around comments if the comment includes white space.

The following example shows rain gage data for 5 sites during a 60-minute time span ending at 22:00 on 10 July 2000. Note that the exact columns and the numerical precision may vary, as does the order of keywords on each line.

```
REM "Example format of an IRIS Rain Gage input file" TIME 200007102200 SPAN 60
CODE 001213 LONLAT 127.312533 38.172512 RRATE 4.5
CODE 000223 LONLAT 127.2223 38.2155 RRATE 2.2
CODE 000095 LONLAT 127.31 38.2283 RFALL 0.083
LONLAT 127.1214 38.1825 QUAL 0 RRATE 1.2 REM "Gage broken"
CODE 000122
CODE 000109 RRATE 1.1 LONLAT 127.2884 38.1277
```

E.4 Configuring **RAIN1** with Raingage Correction

▶ 1. Configure the input **CAPPI** products.
   The **RAIN1** algorithm starts with **CAPPI** products. See 3.5 CAPPI: Constant Altitude Plan Position Indicator (page 39).
   To optimize the **CAPPI** products for **RAIN1** with **GAGE** correction, consider the following:

   - **CAPPI** data type set to R based on Z.
     The input **CAPPI** should be made in rainfall rate units. This causes the **CAPPI** interpolation to be done on rainfall rate rather than on dBZ. Thus the weighting of the interpolation is linear in rainfall rate units.
     Use the corrected reflectivity factor Z instead of the uncorrected reflectivity factor T so that you can benefit from clutter correction and cancelation.
   - **CAPPI** resolution of approximately 1 km (0.6 mi)
     The resolution of the **RAIN1** product is identical to that of the input **CAPPI** product. Vaisala recommends that the maximum range for the **CAPPI** and the number of pixels be set to yield a resolution of approximately 1 km (0.6 mi). For example, for a 240 maximum range, a 480 × 480 product would have 1 km (0.6 mi) resolution.
   - **CAPPI** height of approximately 1 km (0.6 mi), **Fill** set to **Yes** (**Pseudo CAPPI**)
     Select a single-level **CAPPI** height of approximately 1 km depending on the height of the surrounding terrain. Note that **CAPPI** heights are defined as height above a reference level that is usually set to sea level (in IRIS Setup/RCP). Enable the **CAPPI Fill** so that the **CAPPI** is defined for near and far ranges where it is not possible to interpolate to 1 km.
   - Use default Z−R relationship of $Z=200R^{1.6}$
     Vaisala recommends that you initially use a default $Z=200R^{1.6}$ as the default Z−R relationship in the **CAPPI** product. After some experience, you may choose to modify this to match their precipitation more closely. This relationship can be adjusted seasonally as well.

   2. Configure the **RAIN1** product.
   For the **RAIN1** product configuration select the input **CAPPI**s configured above. Use the same Z−R relationship in the **RAIN1** product as was used for the **CAPPI**s.
   If you want to use gage correction, select the gage correction button.

3. Update the **GAGE** correction configuration file (*Rain1GageCor.conf*)
   The raingage correction algorithm uses the *Rain1GageCor.conf* configuration file in the */usr/sigmet/config* directory.
   Copy the template in the */usr/sigmet/config_template* directory to */usr/sigmet/config* during installation.
   The following is an example *Rain1GageCor.conf* file.

```
# Rain1GageCor.conf for the Rain1/GAGE correction program
LOG NONE
# OVERRIDE_PRODUCT_NAME JOE
# VERBOSE
MIN_VALID_GAGE 1.00
MIN_VALID_RADAR 1.00
MIN_CORRECTION 0.20
MAX_CORRECTION 5.00
```

- `LOG`
  Configures where logging information goes to. Choices are `NONE`, `TERM`, `FILE`, or `BOTH`. `FILE` means place in the *${IRIS_LOG}/Rain1GageCor.log* file.
- `MAX_CORRECTION` and `MIN_CORRECTION`
- The maximum and minimum gage correction factors prevent large, possibly erroneous corrections from contaminating the data.
  The numbers are entered as factors such as 3 and 0.3 for the max and min respectively.
  If the limit is exceeded, the algorithm uses the value of the limit, that is, the correction factors are saturated at the limit.
- `MIN_VALID_GAGE`
  The minimum amount of rainfall hourly accumulation for a gage to be considered as having a valid amount of rainfall.
  This depends to some extent on the accuracy of the gages. A value of 3 mm (0.12 in) would be typical. Note that this corresponds to approximately a one hour average reflectivity factor of 21 dBZ.
- `MIN_VALID_RADAR`
  The minimum **RAIN1** accumulation, averaged about a gage that is valid.
  Radar precipitation accumulations less than this are considered invalid and no gage correction is calculated. This is similar to the Minimum Valid Gage Accumulation, except that in this case it is for the radar. A value of 3 mm (0.12 in) would be typical.
- `MIN_VALID_COVERAGE`
  When the radar **RAIN1** pixels about a gage are averaged (1, 9, 25, 49 pixels, and similar) to make a gage correction, not all of the pixels are generally valid.
  If the percentage of valid pixels is less than this value, than a gage correction is not calculated. For example, if the Minimum Percentage is set to 50 %, and averaging about a gage is over 3 × 3 pixels, then at least 5 pixels must be valid for a gage correction to be generated.
- `OVERRIDE_PRODUCT_NAME`
  Use this command if you wish to change the product name of the **RAIN1** product.
- `RADAR_AVERAGE`

This is the distance in km around the raingage where the radar reflectivities shall be averaged. In the case where the input **CAPPI** resolution is 1 km (0.6 mi) (for example, maximum range of 240 km (149.1 mi) and 480 × 480 pixels), setting the Averaging Distance to 1.5 km (0.9 mi) would result in averaging over 3 × 3 pixels centered about the pixel containing the gage. The following figure shows some examples.

Radar Data Averaging about a Raingage Location
(Input CAPPI pixel resolution set to 1 km for example)

Averaging Distance Set to 0.5 km
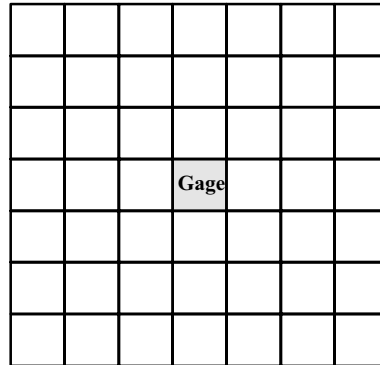
Averaging Distance Set to 1.5 km

Averaging Distance Set to 2.5 km

4. Define IRIS **Setup** configurations.

   a. In **Setup > Input**, specify the input characteristics of the user raingage files including the directory where the files arrive and the name of the Input Pipe that is used to reformat the data in the IRIS **GAGE** product.
   See *IRIS and RDA Utilities Guide* .
   Depending on your system there may also be a `.conf` file for the Input Pipe. The input pipe is stored in `/usr/sigmet/config/pipes`.

   b. In **Setup > Product**, configure the default Z–R relationship.
   Typically this is set to $Z=200R^{1.6}$ (input 200 and 1.6). However, you can set it differently based on local experience.
   Note that you can override this in the **RAIN1** product configuration.

   c. In **Setup > Product**, configure the **GAGE** Wait Time.
   This is the time required after the end of the hour for the **GAGE** product to be generated. The delay is primarily caused by the delay in the arrival of the user raingage data over the network.

E.5 Algorithm for **RAIN1** Gage Calibration

The default raingage correction algorithm provided is modeled after Brandes and Wilson (1982). The algorithm uses a basic time step of one hour, that is, **GAGE** totals for each hour are compared to the hourly radar totals (based on a default Z–R relationship) for the same hour. The ratio (1-hour gage accumulation/1-hour radar accumulation) is the correction factor for a gage.

The **RAIN1** product is then corrected by interpolating the correction factors for all of the gages to each pixel in the **RAIN1** product and multiplying the radar pixel value by the correction factor. The basic algorithm steps are as follows:

1. Gage Correction Factor Calculation for Each Gage
   The gage correction is the ratio of the (1-hour gage accumulation)/(1-hour radar accumulation).
   • The gage accumulation is taken from the **GAGE** product. It must be greater than the Minimum Valid Gage Accumulation for a gage correction factor to be calculated.
   • The radar accumulation for the gage is calculated by first making an uncorrected **RAIN1** product for the hour and then averaging the pixels surrounding the gage (according to the Averaging Distance - see E.4 Configuring RAIN1 with Raingage Correction (page 382)).
   The pixels are averaged by summing only the valid (non-thresholded) pixels and then dividing by the number of valid pixels.
   For a gage correction to be calculated, the average rainfall accumulation must be greater than the Minimum Valid Radar Accumulation, and the number of valid (non-thresholded) pixels must be greater than the Minimum Percentage Radar **RAIN1** Coverage.
   • If both the gage and radar 1-hour accumulations are valid, then a gage correction is calculated for the gage. Otherwise, the gage is flagged as invalid.
   • The maximum and minimum correction factor limits are then applied.
   • This procedure is repeated for each gage.

2. Gage Correction Interpolation for Each Pixel
   The gage correction factors are then applied to the uncorrected **RAIN1** product to yield the final corrected **RAIN1** output product. The technique is that the rainfall accumulation value at each pixel in the uncorrected **RAIN1** product is multiplied by an average correction factor.
   The average correction factor for a pixel is calculated from all gages by weighting the correction factor for each gage inversely with its distance from the pixel, that is,

$$\bar{C} + \frac{\Sigma\frac{C_i}{D_i}}{\Sigma\frac{1}{D_i}}$$

where $\bar{C}$ is the average gage correction for a pixel, C_i is the correction at a gage and D_i is the distance between the gage and the pixel. Note that the minimum value of D_i is fixed at 1 km (0.6 mi).

E.6 **RAIN1** Scheduling with **GAGE** Product

The scheduling algorithm for generation of the **RAIN1** product is different from the standard **RAIN1** scheduling since the product generator must await the arrival of both the **CAPPI**s and the **GAGE** product for the previous hour. To account for the delay in the arrival of the **GAGE** product, a **Gage Wait Time** is defined in IRIS **Setup > Product**. In the following figures, this is labelled `WaitG`.

The notation shows a time line which represents the data time, that is, in the case of a **CAPPI**, this is the time that the volume scan was started based on the clock on the radar computer which may be different than the clock on the computer where **RAIN1** is generated. In the case of a **GAGE** product this is the time exactly on the hour, at the end of the accumulation hour. Vertical bars show even hours.

Case 1A: Normal Case, CAPPIs and GAGE Arrive

This is the normal case when both the **CAPPI** and the **GAGE** products arrive. The **GAGE** product arrives sometime after the hour, but in `Data Time` it is plotted exactly on the hour. The **RAIN1** product is run as soon as the **GAGE** product arrives and the first **CAPPI** from the next hour arrives.

Case 1B: CAPPIs arrive, GAGE Arrives for a Future Hour

Here there is a missing hour in the **CAPPI** product stream. The **RAIN1** product runs (without **GAGE** correction). The assumption here is that the **GAGE** products are arriving in time order, and the missing hour's data never appears.

Case 1C:CAPPIs Arrive, but GAGE Does Not

Here the **GAGE** product does not arrive. The **RAIN1** product runs (without **GAGE** correction) when the local time on the product computer is past the arrival time of the last input **CAPPI** plus the `Tail` and `WaitG` times. The `Tail` time is difference between the data time of the last input and the next integer hour. The `WaitG` time is the `Raingage Arrival Wait Time` from **Setup > Product**.

Note that if the **GAGE** product eventually arrives, it is ignored.

Case 2A: GAGE Arrives, But the First CAPPI in the Next Hour Does Not

This is similar to the previous case. In this case, after the `Tail` and `WaitG` times have elapsed, the product runs using the available **CAPPI**s in the prior hour. There is a **GAGE** correction and the **RAIN1** algorithm uses all the available **CAPPI** products.

Case 2B: Neither GAGE nor the First CAPPI in the Next Hour Arrive

This is nearly identical to the Case 2A, that is, after the `Tail` and `WaitG` times have elapsed, the product runs using the available **CAPPI** products in the prior hour except that there is no **GAGE** correction.

Degenerate Case: No CAPPI's Ever Arrive

In this case, no **RAIN1** product is generated.

More Information

E.7 Hydromet Raingage References

Brandes, E. A. and J. W. Wilson, 1982: Thunderstorms: A social, Scientific and Technological Documentary; Vol 3: Instruments and Techniques for Thunderstorm Observation and Analysis. E. Kessler, Editor. U.S. Department of Commerce, NOAA Environmental Research Laboratories.

Appendix F. Task Configuration Examples

F.1 Task Configurations

You can customize and add to the default task configurations.

Note the following guidelines when creating or modifying tasks:

1. As it is not possible to get the best velocity and reflectivity data at the same time, default tasks are divided as follows:
   - Long-range surveillance with coarse spatial resolution for detecting weak echoes (`SURVEILLANCE` task).
   - Medium-range polarimetric volume with good spatial resolution for high quality reflectivity and polarimetric data (`VOL_A` and `VOL_B` tasks).
   - Short-range volume with good spatial resolution for high quality velocity data (`WIND` task).
2. To maximize signal processing data quality, all tasks use 2-byte data in the radar server computer.
   By default, all RAW products are truncated into a 1-byte form for transmission. If the bandwidth allows, use 2-byte data in RAW products.
3. After installing and calibrating the radar system or if you change the default task processing parameters, you must adjust the quality threshold values.
   For example, if you increase the number of samples, enter a slightly decreased LOG value to improve measurement sensitivity.
4. The tasks use Doppler filter number 4, which is a Gaussian adaptive (GMAP) filter with a 0.4 m/s width.
   Depending on the clutter conditions, test a weaker filter (number 3 or smaller) or a stronger filter (number 5 or greater) to find the optimal filter for your radar site. All the tasks use a CSR threshold of 40 dB.
   Depending on the clutter conditions, you can test a stronger value (less than 40 dB). Before finalizing the setting, measure with real weather to verify that the weather echoes near zero velocity do not vanish (Doppler snake).
5. *Dual pol.*: To benefit from the enhanced sensitivity of the polarimetric signal processing, use the `Ze` data moment to generate echo products.

F.1.1 SURVEILLANCE

The `SURVEILLANCE` task is optimized for long-distance detection of precipitation echo. It uses the longest available pulse width of the weather radar system.

Dual polarization systems use the enhanced reflectivity `Ze` along with a large number of processed samples to significantly improve the sensitivity when compared with single polarization systems. Samples are collected from a large contributing area using a wider azimuthal spacing of the radar rays (2°) and a range averaging over several consecutive range bins.

Ifffffffffff

The default configuration has a maximum range of 450 km. Use it in warm climates. In cooler climates, lower the maximum range to, for example, 350 km in summer or 250 km in winter. Shorter ranges allow for higher pulse repetition frequencies (within the range) and more pulses can be configured within the constant run time.

The LOG threshold is optimized for tropical conditions with a clear margin of noise uncertainty, even in conditions of extreme rain. In cooler climates, with a calibrated radar, the LOG threshold can be as low as 1.2 dB.

The following figure shows the task configuration details for a dual polarization system. Blue circles indicate items with different values in single polarization systems.

Figure 127   SURVEILLANCE Task Configuration -Dual Polarization

Azimuth Resolution	1.00°
Polarization	Horiz
Samples	35
DP Attn Cor Z ZDR	disabled
Range Avg/Smt	2
Proc Mode	FFT
Measured data moments	T, Z, Vel, Width, SQI

F.1.2 VOL_A

The VOL_A task is the lower part of the hybrid volume scan VOL_A and VOL_B. It is optimized for precipitation observations and measurements within a range of 250 km. Use a single pulse repetition frequency mode for the maximum quality of Z and dual polarization moments.

Dual polarization systems use the enhanced reflectivity `Ze` to improve sensitivity. The Z data fields are quality-controlled for maximum rejection of non-meteorological echo. The Z data presents the true fields of precipitation. The combined observations of KDP and Z are used in quantitative rainfall estimates.

In difficult conditions with moving clutter, sea clutter, and interference, you can raise the PMI threshold to as high as 0.55. Try a higher SQI value with some loss of precipitation in the melting layer and in strong convection.

The following figure shows the configuration details of the task for a dual polarization system. Blue circles indicate items with different values in single polarization systems.

Figure 128   VOL_A Task Configuration - Dual Polarization

Polarization   Horiz
DP Attn Cor Z ZDR   disabled
Proc Mode   FFT
Measured data moments   T, Z, Vel, Width, SQI

F.1.3 VOL_B

The VOL_B task fills the higher elevations of the hybrid volume providing observations of standard and dual polarization moments for volume products, such as CAPPI.

The task uses high pulse repetition frequency and delivers high quality radial winds in the range of +/- 16 m/s. Due to a high pulse repetition frequency, high elevations, and the modest ranges in use, narrow precipitation spectra are possible. The quality of standard and dual- polarization moments is good. The hybrid volume data of VOL_B offer inputs for a variety of meteorological products.

At higher elevation angles:

- The system is less likely to detect second trip echoes or non-meteorological targets and thus, SQI and PMI thresholding is not used for Z data fields by default, unlike in VOL_A.

- Enhanced reflectivity **Ze** is not selected like for short measurement ranges of high elevations as it does not bring significant sensitivity improvements.

The following figure shows the configuration details of the task for a dual polarization system. Blue circles indicate items with different values in single polarization systems.

Figure 129   VOL_B Task Configuration - Dual Polarization

Polarization	Horiz
DP Attn Cor Z ZDR	disabled
Proc Mode	FFT
Measured data moments	T, Z, Vel, Width, SQI

F.1.4 WIND

The WIND task is a short distance volume scan providing wind observations for products such as VVP and WIND.

The task uses the FFT major mode in H-only. A robust dual pulse repetition frequency provides unambiguous winds of up to +/- 50 m/s, up to a range of 78 km. The quality settings are those of a classic Doppler.

A higher dual pulse repetition frequency ratio (4:3, 5:4) can increase the unambiguous velocity regime and/or the maximum range. However, using higher dual PRF ratios as well as longer ranges decreases the quality of velocity data.

Figure 130   WIND Task Configuration

F.1.5 SUNCAL and ZDRCAL Tasks

Detailed configurations of the SUNCAL and ZDRCAL tasks are in the */usr/sigmet/config/suncal.conf* and *zdrcal.conf* files.
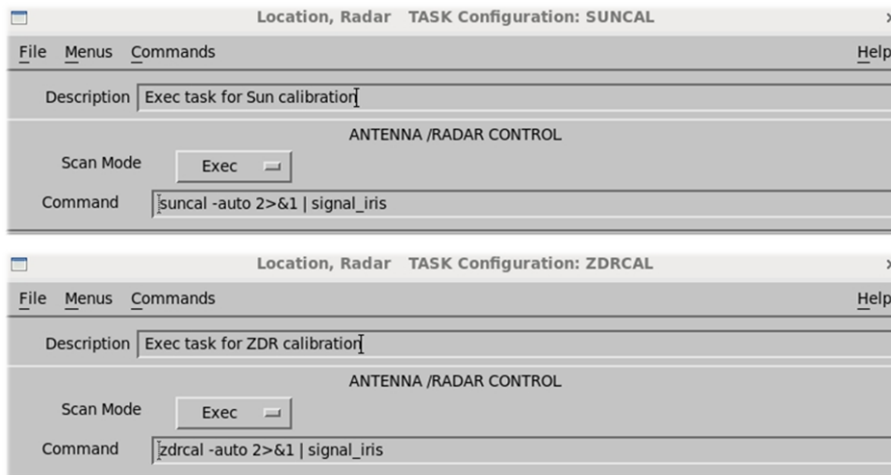
Figure 131   SUNCAL and ZDRCAL Task Configurations

SUNCAL

For the SUNCAL task, adjust the start, stop, and repeat time for each site so that the task is executed three times a day: shortly after sunrise, at noon, and just before sunset.

ZDRCAL

ZDRCAL task is only for dual polarization systems.

F.2 Product Configurations and Scheduler

F.2.1 RAW Products

The radar system produces **RAW** products on the radar server computer and sends them to the central site for product generation.

To reduce communication bandwidth, by default data is truncated to1- byte, 8-bit resolution from the original 2-byte, 16-bit resolution.

If the communication bandwidth allows, use 16 bit resolution to improve data quality for post-processing:

▶ 1. On the product configuration menu, select **Type > RAW**.
By default, the option is **Forced 8-bit Format**, which truncates the content.

2. Select **Preserve Format** to preserve 16-bit resolution.
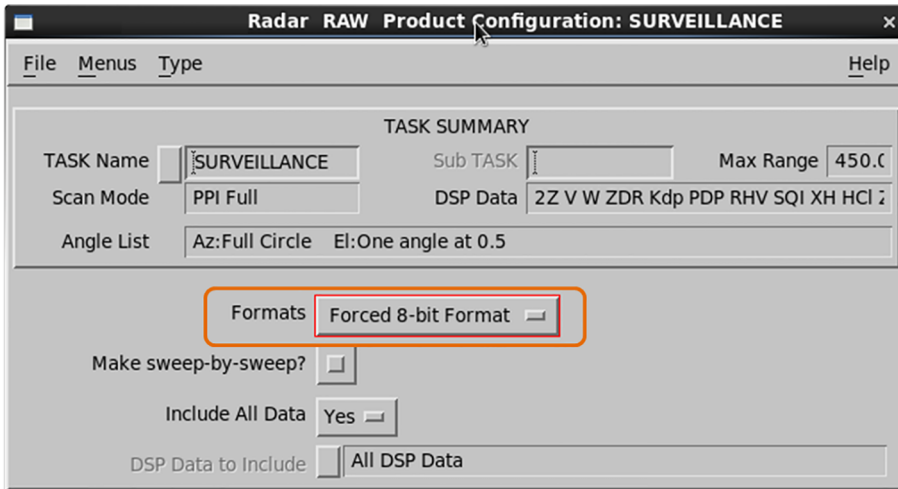
3. Select **File > Save As > OK**.

F.2.2 Melting Height Products

In dual polarization systems, a melting height product, VOL, is configured and scheduled by default.

This product uses tasks VOL_A and VOL_B to compute the height of the melting layer.
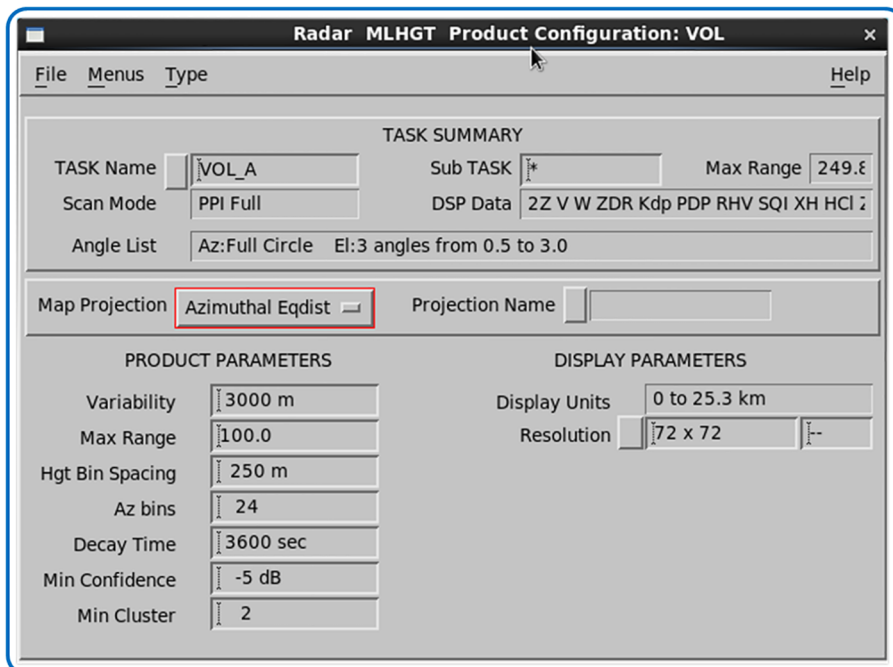
Figure 132   Melting Height Product Configuration

F.2.3 Product Scheduler

If you need other products at the radar site, add them to the **Product Scheduler**. This
applies to both radar site and central site data production.

Input Task(s)	Products
Surveillance	Weather reflectivity surveillance.
VOL_A and VOL_B	Reflectivity-related end user products, such as rain or storm detection or hydrological measurements.
WIND	Wind-related products. For example, V, VVP, WIND, and SHEAR.

1) *In dual-polarization systems, use the R (Z, KDP) algorithm for QPE products, such as rainfall rate or accumulation.*
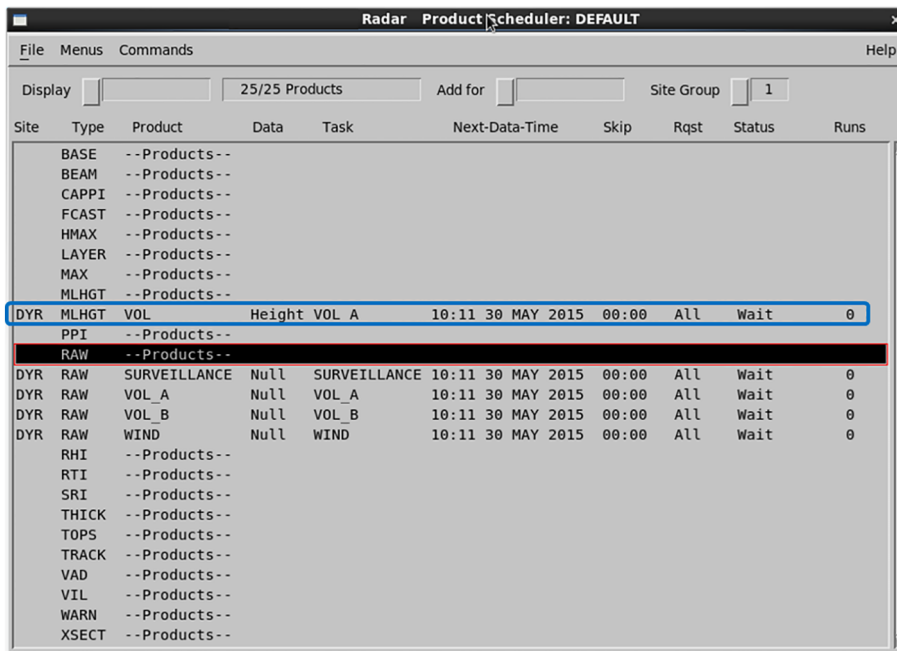
Figure 133   Default Product Scheduler

Index

Warranty

For standard warranty terms and conditions, see www.vaisala.com/warranty.

Please observe that any such warranty may not be valid in case of damage due to normal wear and tear, exceptional operating conditions, negligent handling or installation, or unauthorized modifications. Please see the applicable supply contract or Conditions of Sale for details of the warranty for each product.

Technical Support

Contact Vaisala technical support at helpdesk@vaisala.com. Provide at least the following supporting information:

- Product name, model, and serial number
- Name and location of the installation site
- Name and contact information of a technical person who can provide further information on the problem

For more information, see www.vaisala.com/support.

Recycling

Recycle all applicable material.

Follow the statutory regulations for disposing of the product and packaging.

VAISALA