

## 9.3 AIRDAT SYSTEM FOR ENSURING TAMDAR DATA QUALITY

Alan K. Anderson \*  
AirDat, Evergreen, Colorado

### 1 ABSTRACT

AirDat is equipping commercial aircraft with a network of TAMDAR (Tropospheric Airborne Meteorological Data Reporting) sensors. The network uses communication satellites to relay atmospheric observations in near real time to a data processing center. It is important to maximize the amount of high quality atmospheric data, and minimize the amount of bad data, before it is delivered for use by operational meteorologists, aviation support systems, and forecasting models. Questionable data, or a decrease in the reliability of a sensor, must be quickly identified and acted upon. AirDat developed a quality assessment system to help achieve these goals. The system structure permits rapid automated responses to suspicious data and facilitates the flexible insight and problem identification skills that people provide.

Characteristics of the AirDat system used to monitor the quality of high volumes of meteorological data in a timely manner are described in this paper. It explores the type of problems detected by automated systems, the role of meteorologists and engineers in the quality assessment process, and how the system will evolve to incorporate faster automated responses to suspicious data.

Three key elements in the quality assurance process are considered:

1. "Real-Time Quality Filters" perform immediate assessment of incoming observations and determine what measurements are appropriate for distribution.
2. A separate system – named "Delta Hound" – performs an exhaustive automated analysis of observations. It looks for suspicious patterns, compares TAMDAR measurements to information available from reference sources, and provides tools to view the analysis results and baseline data.

3. "AirDat Personnel" use Delta Hound to evaluate the health of TAMDAR sensors and investigate data quality issues. They make high-level decisions about sensor reliability, data filtering, and repairs based on the Delta Hound analysis. Active human participation in the quality assessment process also provides the insight required for new problem detection algorithms and future enhancement of automated systems. Algorithms are initially implemented in Delta Hound; and, upon proving their worth, may form the basis of new real-time quality filters.

### 2 BACKGROUND

AirDat developed a multi-function atmospheric sensor called TAMDAR (Tropospheric Airborne Meteorological Data Reporting). Commercial aircraft operating in North America are being equipped with TAMDAR units to measure humidity, pressure, temperature, winds aloft, icing, turbulence, location, time, and altitude. These measurements are relayed via Iridium communication satellites to the AirDat data center where they are processed, evaluated, formatted, and distributed in near real time.

Increasing the density of atmospheric measurements is a primary goal of the AirDat network of TAMDAR sensors; this can improve mesoscale modeling and aviation safety. Only if the measurements are reliable and accurate will they provide value to meteorologists and forecasting systems. The AirDat quality assurance system, described herein, examines atmospheric measurements and attempts to quickly identify and respond to any inaccuracies.

### 3 SYSTEM OVERVIEW

The following diagram illustrates the relationship between the four key parts of the quality assurance system.

---

\* Corresponding author address: Alan K. Anderson, AirDat, 30726 Bryant Drive Suite #410, Evergreen, CO 80439 U.S.A.; e-mail address: aanderson@airdat.com.

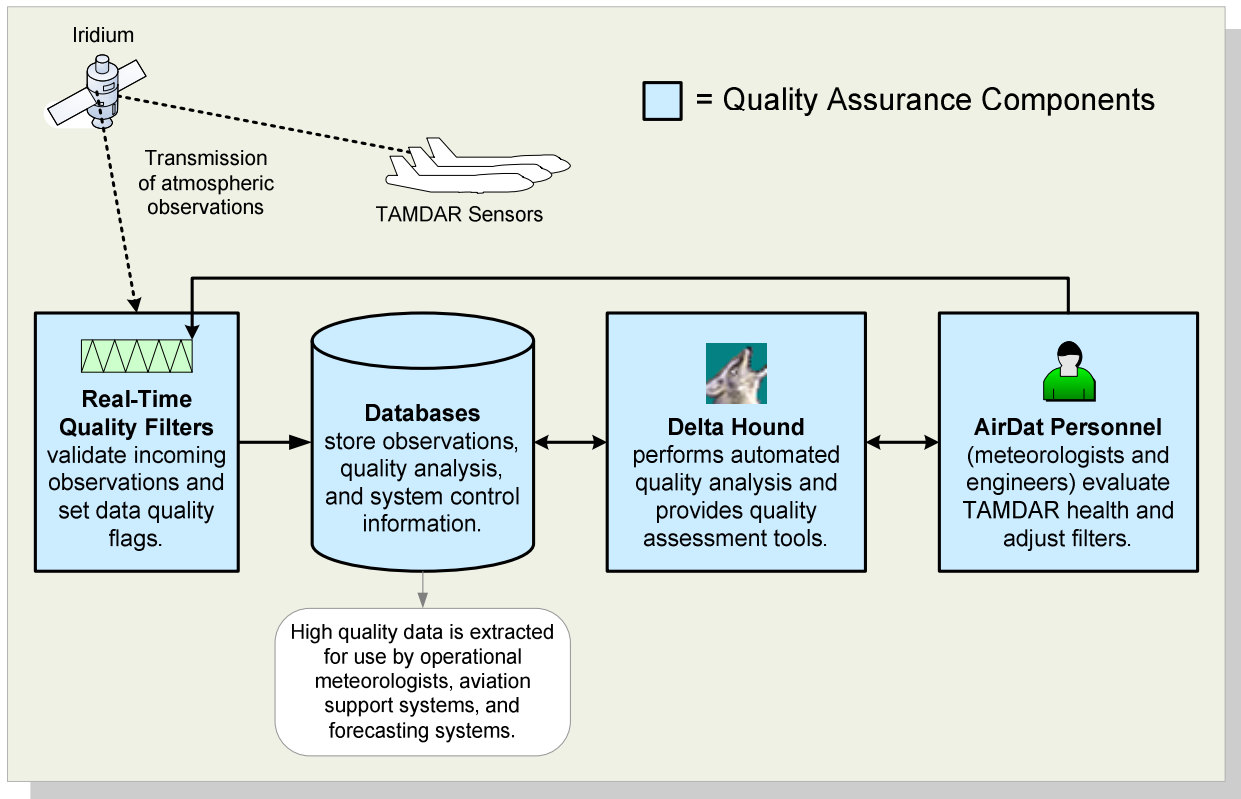


Figure 1: TAMDAR system overview, with an emphasis on quality monitoring system components.

Raw data from a TAMDAR unit typically arrives at AirDat 1-5 minutes after measurements are made. The values within an observation (such as temperature, aircraft speed, etc) are evaluated and immediately assessed by “Real-Time Quality Filters” and quality information is assigned to each value.

The atmospheric observations and associated quality tags are immediately deposited into a database for storage and distribution. Measurements tagged as “good” can easily be extracted from the database for use by operational meteorologists, aviation support systems, and forecasting systems. AirDat quality monitoring is a first step; it is not intended to reduce the need for modeling systems to include a separate assessment of the value of atmospheric measurements used as input.

The heart of the AirDat quality assurance system is known as “Delta Hound”. This sentry is continually on guard, sniffing through the observations and looking for suspicious patterns. It also compares TAMDAR measurements to predictions from forecasting system. Additionally, Delta Hound provides important tools to view the analysis results as well as the baseline measurements.

“AirDat Personnel” use Delta Hound to monitor and investigate the health of TAMDAR sensors. They issue maintenance orders and adjust the real-time filters to accept or reject specific measurements from specific TAMDAR units.

The next section provides a simple example. Then, subsequent sections provide detail about each stage of the quality monitoring process.

## 4 QUALITY ASSESSMENT

### EXAMPLE 1

Each TAMDAR unit is equipped with two relative humidity sensors (known as RH1 and RH2). This section describes a simple quality assessment scenario in which the relative humidity sensors disagree.

In observation A, a normally healthy TAMDAR unit reports a humidity detected by RH1 is 41% and RH2 is 43%. The real-time filters let these measurements pass and declare the “consensus” relative humidity to be 42% and assign an uncertainty value of  $\pm 3\%$  (indicating high quality data).

In observation B, RH1 reports 120% and RH2 is 90%. The real-time filters decide that since RH1 is outside of the sensible range it should not be trusted, the 90% value reported by RH2 is used as “consensus” relative humidity and a larger uncertainty value (indicating lower quality data) is assigned to the measurement.

If a few out-of-range measurements are received from RH1, Delta Hound notes the fact that RH1 and RH2 occasionally disagree, RH1 is the likely culprit, and assigns a lower overall health score for the TAMDAR sensor. If a significant number of out-of-range measurements are received from RH1, Delta Hound issues an alert about the suspicious data.

The person receiving the alert investigates and decides upon an appropriate response. If RH1 is frequently reading too high, typically the real-time filters are adjusted to disregard any measurement from RH1 from that specific TAMDAR unit. Delta Hound stops issuing the alert, but will continue to evaluate the data coming in from the both relative humidity sensors.

From this point, three scenarios are possible.

1. If Delta Hound sees that final relative humidity measurements (from RH2, the remaining good sensor) are in range and typically agree with weather model humidity predictions, then AirDat continues to operate the unit using a single humidity sensor.
2. If Delta Hound sees that the humidity measurements from RH2 (the remaining good sensor) do not agree with model predictions, AirDat will filter out all humidity information from the TAMDAR unit, issue a repair order, and an aircraft mechanics will replace the humidity sensor circuit board.
3. If the problem with RH1 resolves itself and starts agreeing with RH2 – Delta Hound issues an alert indicating that RH1 has recovered, and AirDat personnel may adjust the real-time filters to permit utilization of humidity information from RH1.

## 5 REAL-TIME QUALITY FILTERS

When transmitted atmospheric observations arrive at the AirDat data center, they are immediately evaluated. Using numerous tests – each temperature, wind, humidity, icing, and turbulence measurement is assigned a quality flag that determines what measurements are appropriate for distribution. This ground-based filtering process is able to consider a variety of factors such as the aircraft type and the current status of various sensor components.

The real-time quality filters can be divided into three stages (as depicted in the following figure):

1. Quarantine filters
2. Measurement assessment filters
3. Interdependency checks

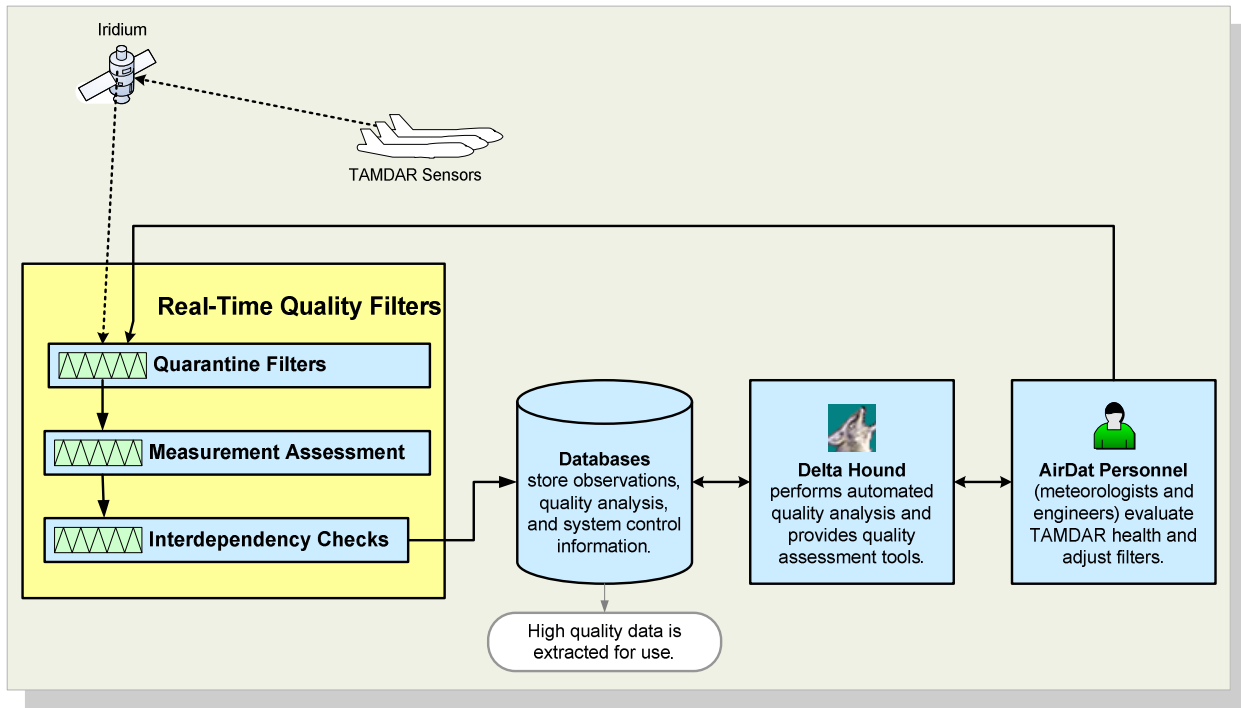


Figure 2: Three stages of the real-time quality filters.

**Quarantine filters** allow AirDat personnel to mark various types of measurements from specific TAMDAR sensors as reliable or questionable. For example, when AirDat first installs a TAMDAR unit AirDat quarantines all data from a probe until it proves itself reliable for a couple days; then the overall quarantine filter for the new unit is turned off. If a particular aspect of a TAMDAR unit proves unreliable, AirDat can quarantine a subset of the data (such as just icing or relative humidity information) until the unit is repaired.

In the next stage, **measurement assessment filters** automatically mark individual measurements as good, bad, or questionable. Range testing can provide basic sanity checks on the values delivered for individual measurements. For example, relative humidity should be between 0 and 100% and altitude should be between -500' and 50,000'. Other filters provide more intelligence. As altitude changes, the expected valid temperature range also changes. When we check the airspeed, not only does altitude come into play, but data validation also depends on the type of aircraft.

In the TAMDAR multifunction sensor, there are many interrelationships between measurements. Inaccuracy in one measurement (such as temperature) can affect the accuracy of other measurements (such as relative humidity). **Interdependency checks** ensure that before the measurement itself is considered to be valid, every parameter that effects the particular measurement must also be valid. For example, aircraft bank angle can effect accuracy of the wind speed & wind direction.

Additionally, measurements from the two relative humidity sensors located on each TAMDAR unit are combined to determine the final (or "consensus") relative humidity value and the potential amount of uncertainty of the consensus humidity value.

The following table summarizes five important measurement interdependencies.

Measurement	Other Items Affected
Aircraft heading	Wind speed, wind direction, and aircraft bank angle
Aircraft bank angle	Wind speed & wind direction
Icing and deicing heater	Temperature, relative humidity, wind speed, wind direction and indicated airspeed
Indicated airspeed	Temperature, relative humidity, wind speed & wind direction
Temperature	Relative humidity, wind speed, and wind direction

All measurements and the associated quality information are immediately deposited into a database for storage, distribution, and additional analysis by Delta Hound.

## 6 DELTA HOUND

The next quality assessment component is Delta Hound, this system continually examines incoming data, looking for suspicious TAMDAR measurements and patterns.

Much of the automated scrutiny performed by Delta Hound is time-consuming or involves analyzing multiple observations produced across time spans ranging from a few minutes to a few days. Therefore these analyses are not all suitable for real-time processes; nonetheless, they can detect subtle issues or trends that are worth investigating. Delta Hound periodically performs time-consuming problem detection and updates related statistics; these results are cached in database tables. This caching process allows for faster, on-demand reporting of results.

Key aspects of Delta Hound are depicted in the following diagram.

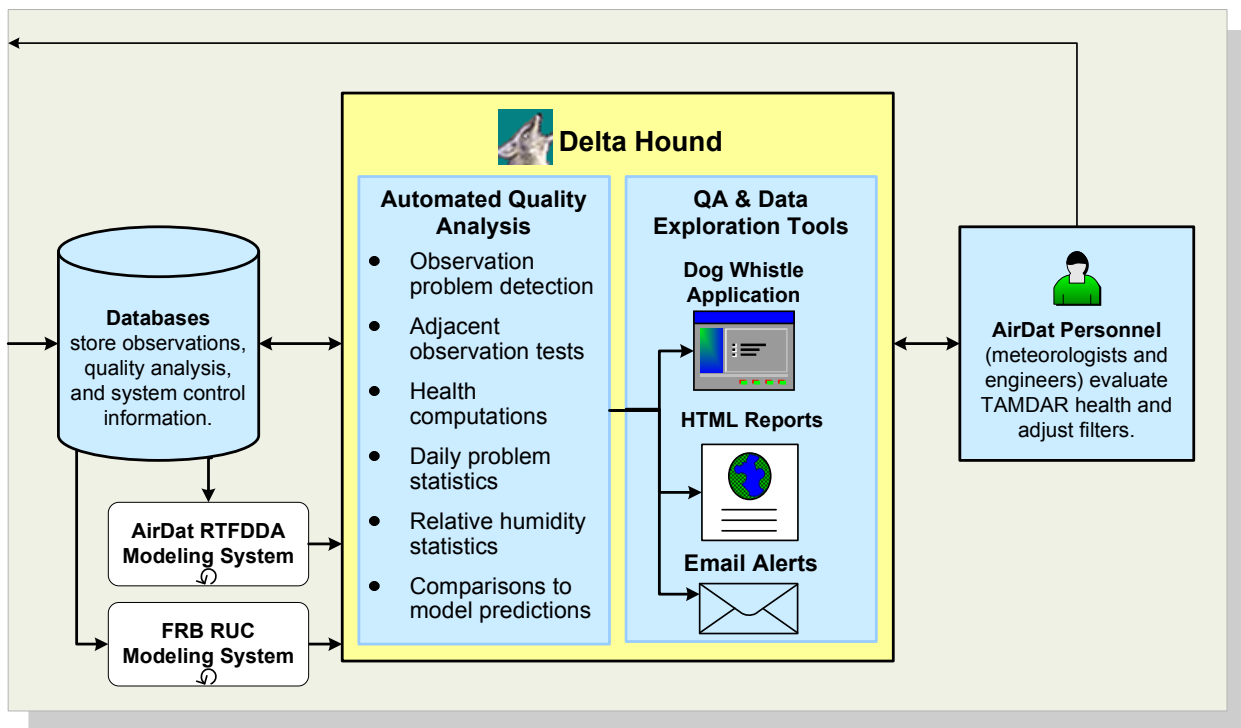


Figure 3: Delta Hound quality analysis routines and data exploration tools.

With 20,000 TAMDAR observations per day – and potentially 10 times that number in the future – it is a challenge to create a system allowing information and analysis results to be investigated in a variety of ways, in a reasonable amount of time. If Delta Hound notices a severe problem it proactively sends out an alert via email. Additionally, it provides data exploration tools for on-demand use. Most information can be viewed using a web browser. A custom Windows application (known as Dog Whistle) is also available to provide high-performance access to the most commonly reviewed information.

The automated analysis of atmospheric measurements performed by Delta Hound can be divided into five categories:

1. Single Observation Problem Detection: identifies problems with individual observations.
2. Adjacent Observation Problem Detection: identifies discrepancies from one observation to the next observation by a single TAMDAR unit.
3. Aggregate Analysis: identifies issues by looking at statistics and patterns across multiple observations.
4. Comparisons to Predictions: identifies discrepancies by comparing atmospheric measurements to forecasting model predictions.
5. Watch Dog Alerts: provide a more human, high-level description of current issues related to TAMDAR data quality.

The following five subsections provide additional information about each automated analysis category:

### **6.1 Single Observation Problem Detection**

Delta Hound examines each observation for the following problems:

- Temperature is unreasonably high or low for the altitude.
- Wind Speed is unreasonably high for the altitude.
- Relative humidity sensors disagree by more than 5% (and neither RH sensor is quarantined).
- Humidity measured by RH1 is outside the valid range (and RH1 is not quarantined).

- Humidity measured by RH1 is outside the valid range (and RH2 is not quarantined).
- Indicated airspeed is outside the valid range for the aircraft and altitude.
- Pressure altitude outside the sensible range or disagrees significantly from the GPS altitude.
- The GPS is not locked or is not responding.
- Airspeed is low (possibly indicating a blocked static pressure port or Pitot tube).
- No aircraft heading data is available.

### **6.2 Adjacent Observation Problem Detection**

Soon after TAMDAR weather observations are added to the database, groups of temporally adjacent observations for each unit are compared for logical inconsistencies. Delta Hound looks for:

- Large time gaps between adjacent observations within a flight.
- TAMDAR rebooting between adjacent observations within a flight (may indicate power problems).
- Rate of change of pressure altitude is outside the expected range (possibly indicating the removal of an obstruction of the static pressure port or Pitot tube).
- Rate of change of temperature is outside the expected range for an ascending or descending aircraft.
- The aircraft is not moving while it is in the air.

### **6.3 Comparisons to Predictions**

AirDat operates a weather modeling system based on MM5 with a Real-Time Four-Dimensional Data Assimilation system (RTFDDA). This generates short-term forecasts for the regions covered by the TAMDAR network. TAMDAR measurements that are deemed “good” by the real-time filters are compared to forecasting model predictions and the deltas (differences) are noted. Deltas for individual observations are rarely noteworthy; however, the statistical measures across many observations are often enlightening.

Similarly, NOAA's Forecast Research Branch (FRB) [formerly Forecast Systems Lab (FSL)] compares TAMDAR data to predictions made by the Rapid Update Cycle (RUC) model. To help to identify problems, Delta Hound also retrieves this data and computes daily statistics of the deltas.

#### **6.4 Aggregate Analysis**

Aggregate analysis seeks to identify problems by looking at statistics across multiple observations. Delta Hound version 1 performs five types of aggregate analysis:

- The frequency of occurrence of each of the (previously discussed) single observation and adjacent observation problems are calculated for each day. (If the frequency of occurrence exceeds a threshold, the "Daily Health and Problems" report uses **red to highlight** the issue.
- Relative humidity statistics are calculated to determine the amount of disagreement between the two humidity sensors and the percent of readings from each sensor that are out of range (24-hour, multi-day, and multi-week statistics are useful for viewing trends).
- Health computations consider all of the problems detected during a day and weights the problems based on the problem severity. The combined problems are compared to the number of observations without problems to generate a single value known as "health". The health number provides a quick numeric summary of the reliability the observations for TAMDAR unit during some time period – usually a day.
- The differences (deltas) between individual TAMDAR measurements and short term forecasting model predictions (RUC & RTFDDA) are computed. The deltas are aggregated to compute daily wind, temperature, and relative humidity bias and error numbers for each TAMDAR unit. It is unwise to draw conclusions based on individual deltas; however, the statistical biases and standard deviations across many observations are often enlightening. When the models are accurate, this information provides a good way to identify an inaccurate TAMDAR sensor.

- Every morning the Delta Hound "Watch Dog" performs a high-level analysis of TAMDAR issues. Looking at the health, observation problems, and statistics for the previous four days – the Watch Dog tries to identify issues that deserve human consideration. This analysis results in a high-level daily investigation report that listing each TAMDAR unit generating a significant number of suspicious measurements and clues about what might be wrong with the unit.

#### **6.5 Watch Dog Alerts**

Every morning the Delta Hound performs a high-level analysis of TAMDAR quality issues (this is called the "Watch Dog"). Looking at the health, observation problems, and statistics for the previous four days – the Watch Dog tries to identify high-level issues that deserve human consideration. These issues include:

- RH1 is often out of range and the RH sensors disagree (neither RH sensor is quarantined).
- RH2 is often out of range and the RH sensors disagree (neither RH sensor is quarantined).
- RH1 is quarantined but humidity measurements look reasonable (RH1 is generally in range and in agreement with RH2.)
- RH2 is quarantined but humidity measurements look reasonable (RH2 is generally in range and in agreement with RH1.)
- Many measurements disagree with RTFDDA model predictions.
- Many measurements disagree with RUC model predictions.
- Suspiciously high temperatures measurements (and temperatures are not quarantined).
- An abnormally high percentage of observations have a specific observation problem.
- The overall health of the TAMDAR unit is good but all observations are quarantined.
- Overall health is poor and observations are not quarantined.
- Overall unit health is declining and observations are not quarantined.

- There has been a long period of inactivity during which an aircraft has not transmitted any observations.
- The ice detector is failing (and icing is not quarantined).
- FDAU (flight data acquisition unit) is not working.
- Shadin converter (that pulls information off the aircraft data bus) is not working.
- Suspiciously high percentage of icing (and icing is not quarantined).
- A TAMDAR unit is transmitting observations, but there is no associated aircraft information.

TAMDAR can also report information about flight departures and arrivals (a.k.a. OOOIs), and the “Watch Dog” detects OOOI problems; however, OOOI-related quality assessment issues are not discussed in this document.

## **6.6 Delta Hound Data Exploration Tools**

In addition to looking for suspicious data, the Delta Hound provides software tools used to monitor TAMDAR data and assess quality.

The tools provide:

- Flexible access to detailed TAMDAR atmospheric observations.
- Forecasting model predictions corresponding to TAMDAR observations and deltas.
- Detailed and summary information about suspicious measurements and the health of TAMDAR units.
- Information about the status of deployed TAMDAR sensors, related equipment, and aircraft flights.
- Historical data and analysis results.

Delta Hound delivers information in three ways:

1. HTML-based reports permit company-wide access using a web browser.
2. A Windows application (known as Dog Whistle) provides high-performance data access for the hard-core personnel who frequently monitor TAMDAR performance and data quality.
3. Email messages alert AirDat personnel of critical TAMDAR problems.

Examples of the Delta Hound tools and reports are provided by the following screen shots.



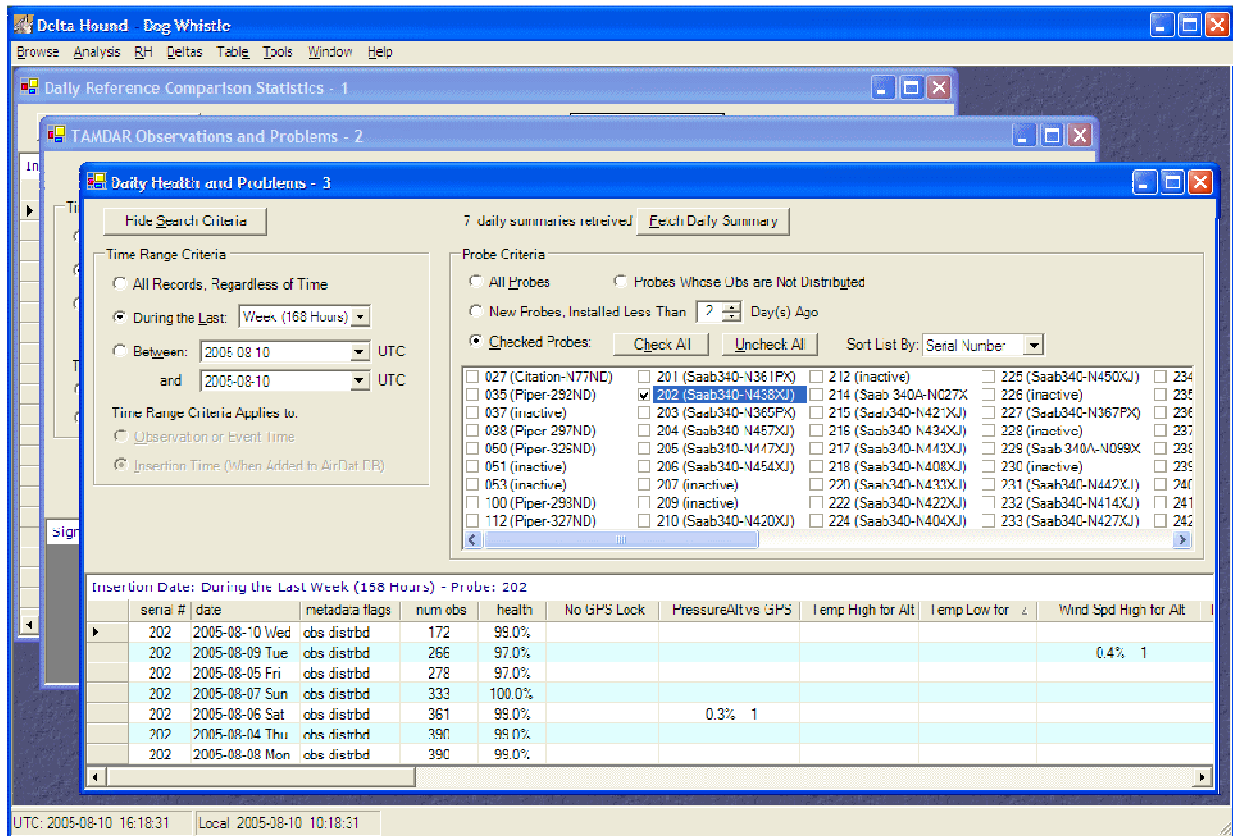


Figure 4: The “Dog Whistle” application. The front-most window shows a typical control panel for specifying what information should be retrieved.

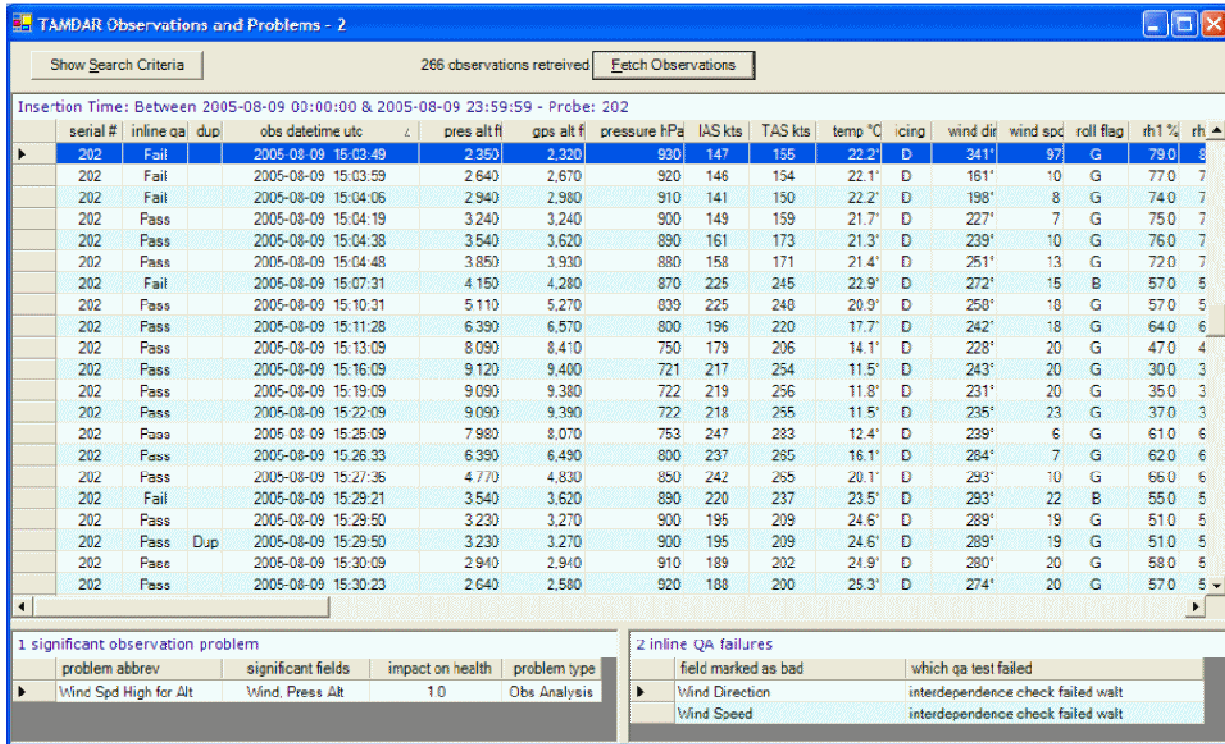


Figure 5: The “Dog Whistle” window for browsing detailed observation data.

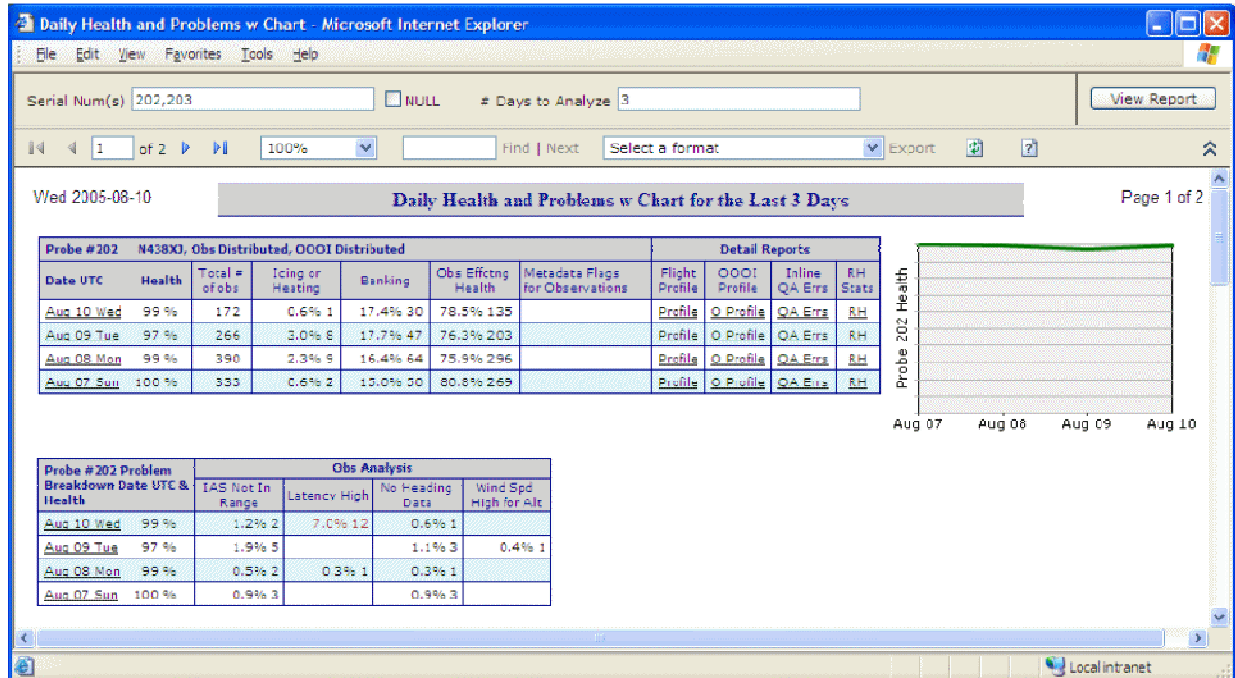


Figure 6: An HTML-style Delta Hound report displayed in a web browser. This report summarizes the activity and problems for a TAMDAR unit during the last few days; it also provides links to drill down and examine more detailed information.

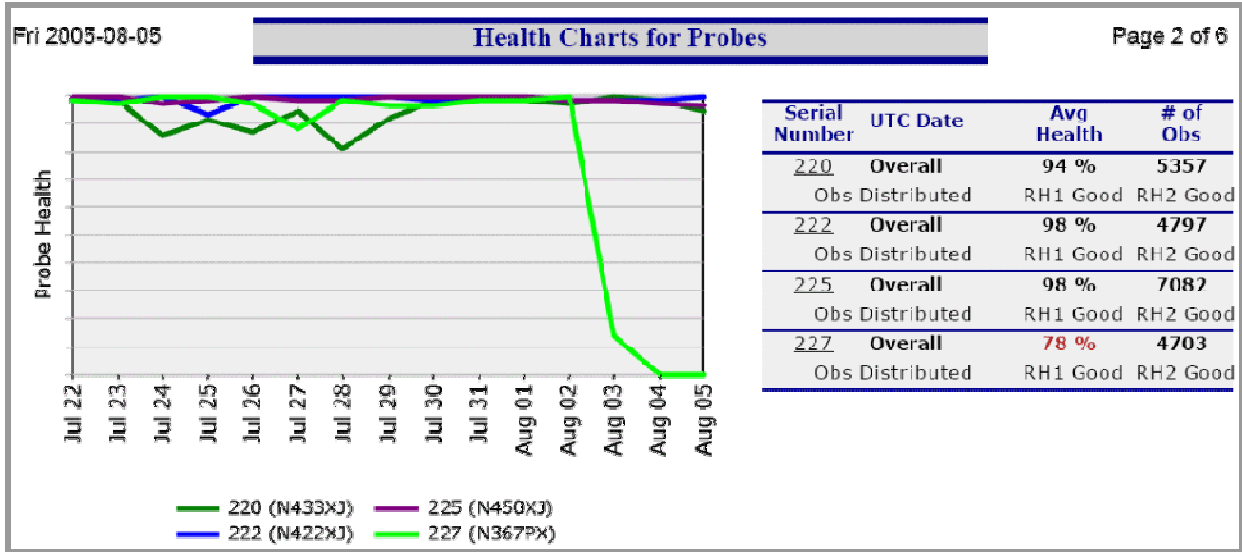


Figure 7: Charts in this report provide a historical overview of the health of TAMDAR units. It is easy to see that something went wrong with TAMDAR serial #227 (or the aircraft with tail number N367PX) on August 3<sup>rd</sup>.

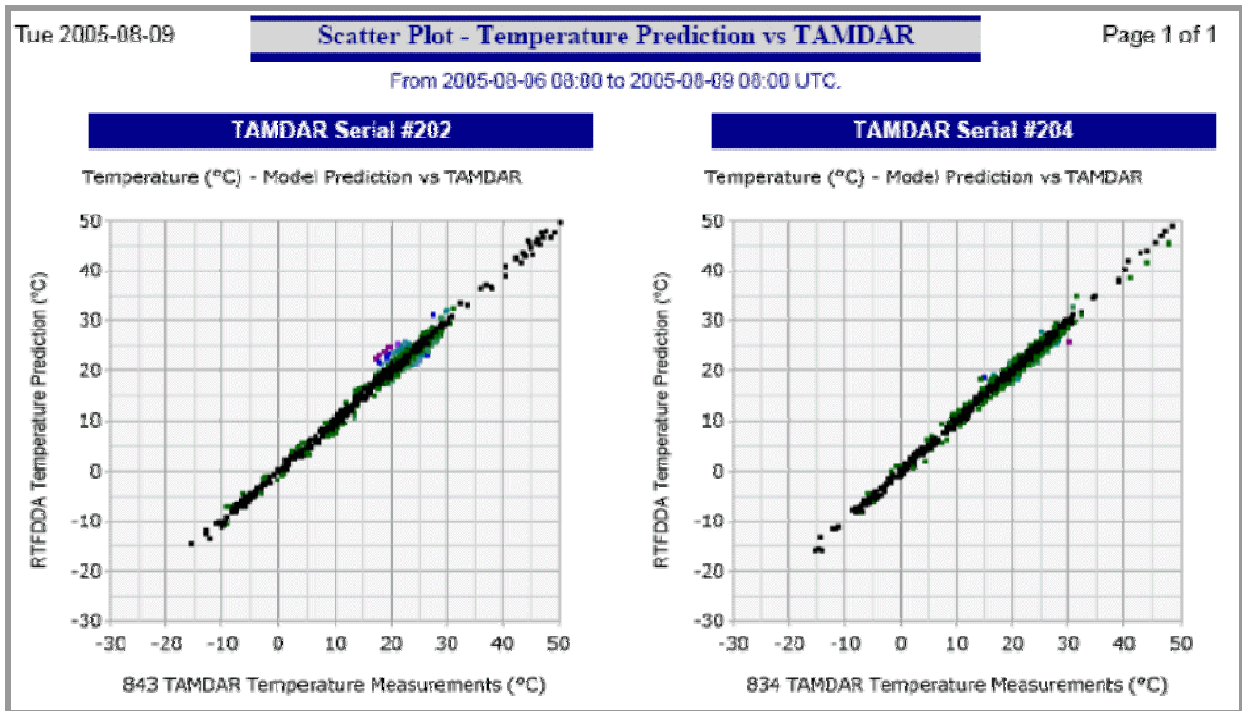


Figure 8: Delta Hound report using scatter plots to show how well temperature measurements for two TAMDAR sensors correspond to temperatures predicted by the RTFDDA. These plots show good agreement between the predicted temperatures and the actual measurements.

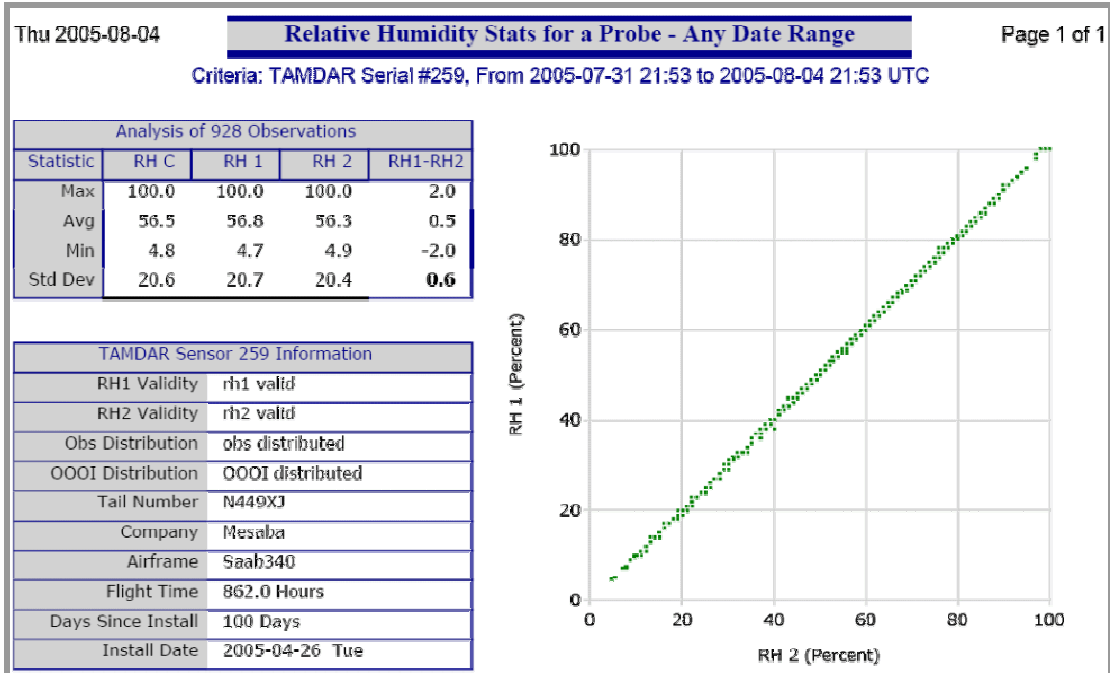


Figure 9: Each TAMDAR unit includes two relative humidity sensors. This report compares measurements from the two sensors (RH1 and RH2). The scatter plot shows good agreement between the two humidity sensors, as does the very low standard deviation of the difference between the two sets of relative humidity measurements. AirDat would have a high confidence in the measurement accuracy from this TAMDAR unit.

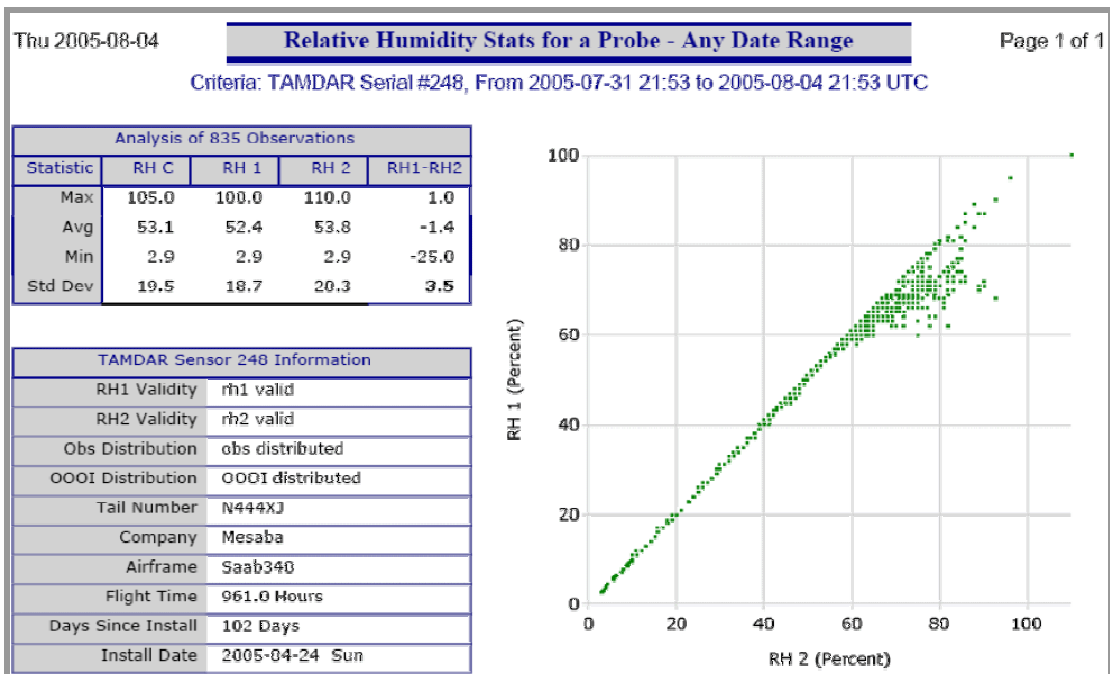


Figure 10: A humidity report showing disagreement between the two humidity sensors; one humidity sensor on this TAMDAR unit is starting to malfunction.

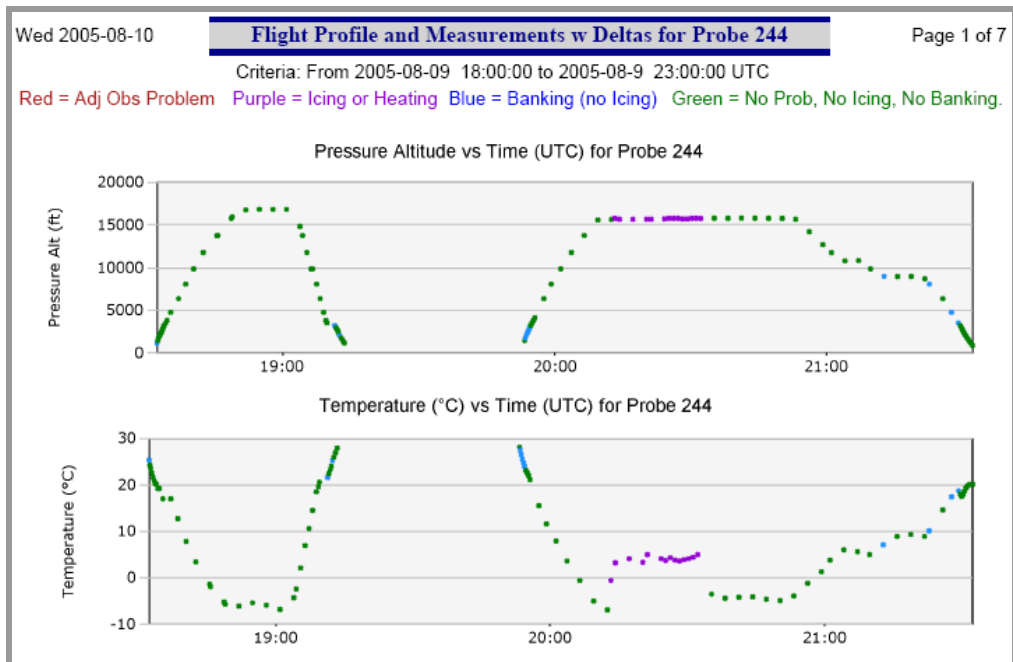


Figure 11: Altitude and temperature charts from a “Flight Profile” report. Plotting measurements as a function of time enhances visualization of flights and the conditions encountered by TAMDAR units. They illustrate the expected decrease in temperature as the aircraft ascends. The purple dots show icing occurred when the second flight leveled off at 15,000’. When icing occurs TAMDAR activates heating elements to melt the accumulated ice. This temporarily elevated the probe temperature to 5°C; the elevated temperatures are marked as invalid by the real-time quality filters’ interdependency checks.

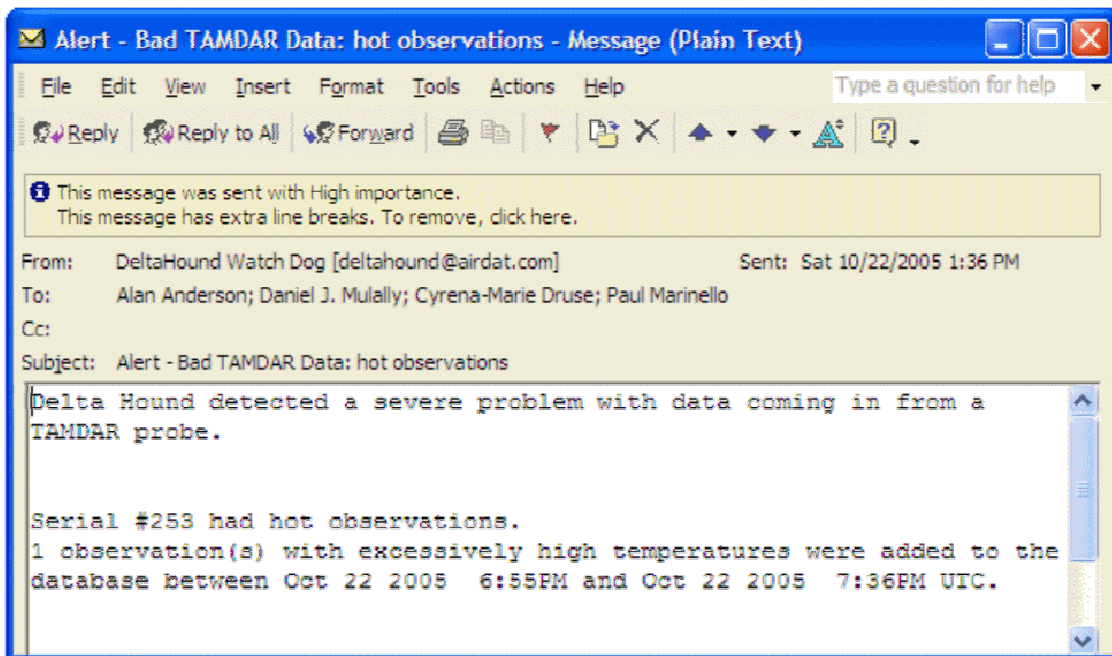


Figure 12: An email alert sent by Delta Hound to AirDat personnel.

## 7 AIRDAT PERSONNEL

In addition to computers, AirDat relies on the judgment of meteorologists and engineers to improve the effectiveness data quality monitoring.

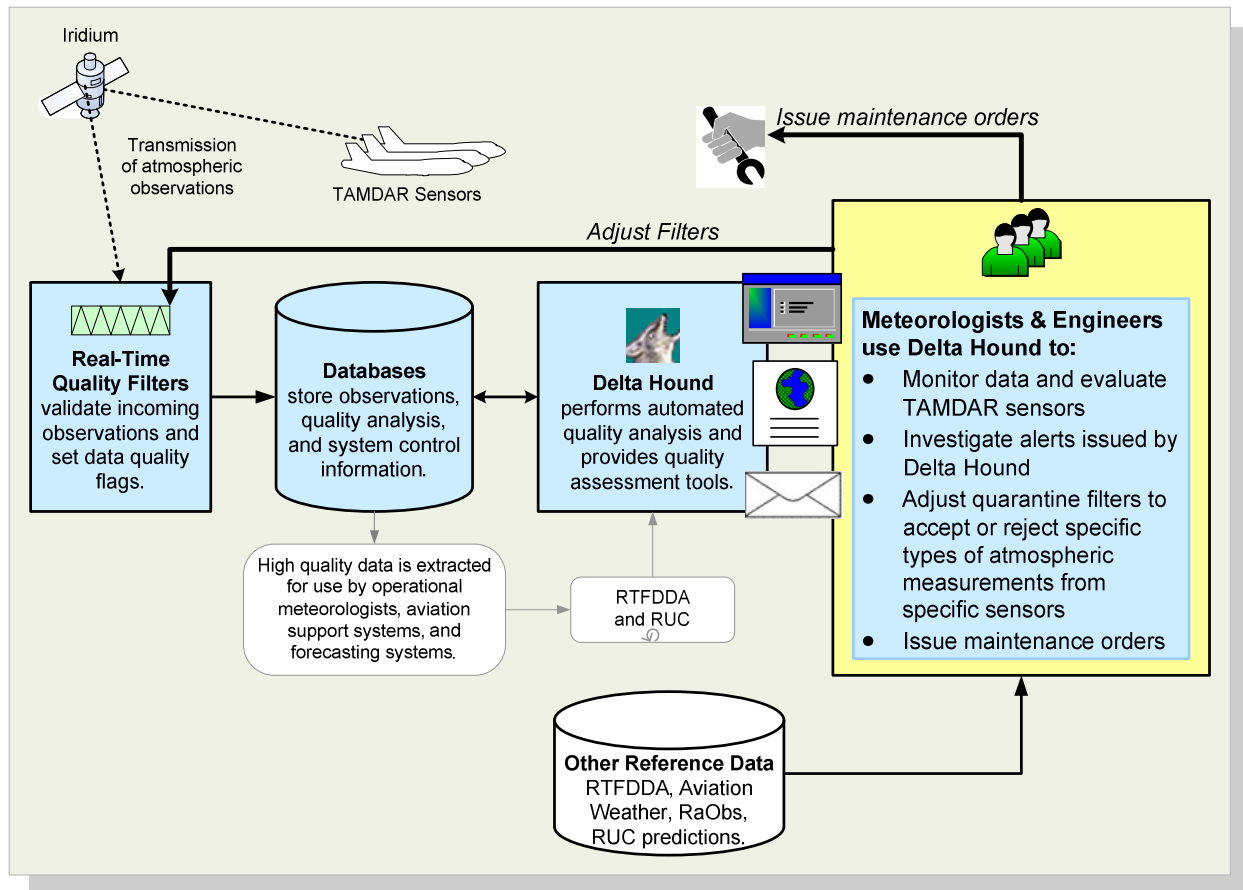


Figure 13: Quality assurance activities performed by AirDat personnel.

Often atmospheric sensor quality is not a “black or white” situation. Human expertise is still the best way to decide how to respond to “gray” situations involving subtle or short-term degradation in measurements. Automated comparison of TAMDAR measurements to other atmospheric information (like model predictions or wind profilers) is a good technique for spotting problems. However, when disagreements exist between TAMDAR and other information, it often requires a skilled meteorologist to determine which data source is exhibiting poor quality.

Including people in the quality monitoring cycle also helps AirDat to spot trends and detect previously unknown problems. Human expertise provides the foundation for developing new algorithms for additional automated problem detection. At AirDat, new algorithms are initially implemented in Delta Hound; and, upon proving their worth, may form the basis of new real-time quality filters.

## 7.1 QUALITY ASSESSMENT EXAMPLE 2

A typical quality assessment session at AirDat begins by looking at the “watch Dog” report – a sample is shown below.

Thu 2005-08-04		Watch Dog - Current QA Alerts			Page 1 of 1
These issues were primarily determined by examining TAMDAR data from 8/1/2005 to 8/3/2005.					
Serial #	Tail Number	Category	Issue Abbrev	Issue Description	
<a href="#">227</a>	328ND	Alert	Poor Health	Health is poor and observations are distributed. (14% health)	
<a href="#">231</a>	N442XJ	Alert	Declining Health	Health is declining and observations are distributed. (04.00% hlt)	
<a href="#">260</a>	N437XJ	Alert	AWOL Aircraft	A long time without observations from a probe assigned to a tail number.	
<a href="#">235</a>	N363PX	Good News	Healing RH1	RH1 is marked as bad, but RH1 is generally in range and in agreement with RH2. (3.19% disagree)	

Figure 14: Part of the Watch Dog report.

Additional details about a particular TAMDAR unit are provided by clicking on the serial number link. The link for serial #227 leads to the “Daily Health and Problems” report shown below. The second table (labeled “Problem Breakdown”) highlights the fact that the poor health relates to an inability for TAMDAR to get heading information from the aircraft’s data bus. The lower table (labeled “Statistical Issues”) points out that on August 3<sup>rd</sup> the Watch Dog thought the FDAU (flight data acquisition unit) stopped functioning.

Fri 2005-08-05		Daily Health and Problems w Chart for the Last 5 Days										Page 1 of 5
Probe #227 N367PX, Obs Distributed, No OOOI Distbn												
Detail Reports												
Date UTC	Health	Total # of obs	Icing or Heating	Ranking	Obs Effctng Health	Metadata Flags for Observations	Flight Profile	OOOI Profile	Inline QA Errs	RH Stats		
Aug 05 Fri	0 %	50	0.0% 0	100.0% 50	100.0% 50	wind bad,	Profile	Q Profile	QA Errs	RH		
Aug 04 Thu	0 %	353	0.3% 1	100.0% 353	95.7% 352	wind bad,	Profile	Q Profile	QA Errs	RH		
Aug 03 Wed	14 %	231	0.9% 2	85.3% 197	96.1% 222		Profile	D Profile	QA Errs	RH		
Aug 02 Tue	99 %	264	0.4% 1	20.0% 55	76.1% 201		Profile	Q Profile	QA Errs	RH		
Aug 01 Mon	98 %	240	0.4% 1	17.5% 42	79.2% 190		Profile	Q Profile	QA Errs	RH		
Jul 31 Sun	98 %	276	0.0% 0	19.6% 54	77.5% 214		Profile	D Profile	QA Errs	RH		
Probe #227 Problem Breakdown												
Date UTC & Health	IAS Not In Range	No Heading Data	PressureAlt vs GPS	Stationary								
Aug 05 Fri	0 %	22.0% 11	100.0% 50	5.0% 3								
Aug 04 Thu	0 %	53.8% 190	100.0% 353	45.3% 160								
Aug 03 Wed	14 %	26.8% 62	83.1% 192	18.6% 43								
Aug 02 Tue	99 %	1.8% 4	1.8% 4									
Aug 01 Mon	98 %	1.3% 3	0.8% 2									
Jul 31 Sun	98 %	1.6% 5	1.1% 3									
Probe #227 Stat Issues												
Date UTC	Poor Health	Bad FDAU	Declining Health	W off Wheels per Flt	W on Wheels per Flt							
Aug 04 Thu	0.00 % hlt			0.00 tOff/flt	0.00 landg/flt							
Aug 03 Wed		83.12 % no heading	0.00 % hlt	0.00 tOff/flt	0.00 landg/flt							

Figure 15: A Delta Hound “Daily Health and Problems” report for a single TAMDAR unit.

Although further investigation could be performed, this situation is clear enough (to someone experienced with TAMDAR) to justify a maintenance order instructing the aircraft mechanics to repair the FDAU.

## 8 CONTINUED DEVELOPMENT

The versatile data displays provided by Delta Hound, combined with the training of AirDat meteorologists and engineers, provides the insight for future enhancement of automated systems. In addition to new problem detection algorithms, AirDat plans to improve TAMDAR quality assessment by:

- Speeding up the RTFDAA forecasting cycles.
- Devising an automated scheme to evaluate the trustworthiness of forecast models. This will help weigh the importance of discrepancies between TAMDAR measurements and the model predictions.
- Letting Delta Hound automatically set some quarantine filters (when automated routines are confident that a significant problem exists).
- Automating comparison of TAMDAR observations to additional references sources – such as RaObs, ACARS, wind profilers, and other nearby TAMDAR observations (a.k.a. Buddy data).
- Tracking changes in sensor calibration over long time periods.

## APPENDIX A: DELTA HOUND ARCHITECTURE OVERVIEW

Delta Hound stores, manipulates, and analyzes large quantities of weather observation data from a variety of sources. The system consists of five primary components (depicted in the following diagram):

1. Data Tier – consists of databases providing storage and flexible access to weather observations, OOOI events, cached analysis results, system control parameters, and aircraft configuration information. Two Microsoft SQL Servers host these relational databases.
2. Middle Tier (a.k.a. the Hound Dog) – performs data comparison and analysis and provides a programming interface for applications to access observation data and analysis results. Although logically separate from the data tier, the middle tier services are implemented using SQL stored procedures and views.
3. Reporting Services and the Web Server (a.k.a. the Groomer) – provide a centralized mechanism creating, managing, and generating nicely formatted reports. A web server delivers the reports to Dog Whistle or web browsers.
4. Application Tier (a.k.a. Dog Whistle) – provides a user interface to display results and let users see the analysis summaries and probe into underlying data.
5. Observation Importers (a.k.a. Retrievers) – fetch reference weather data.



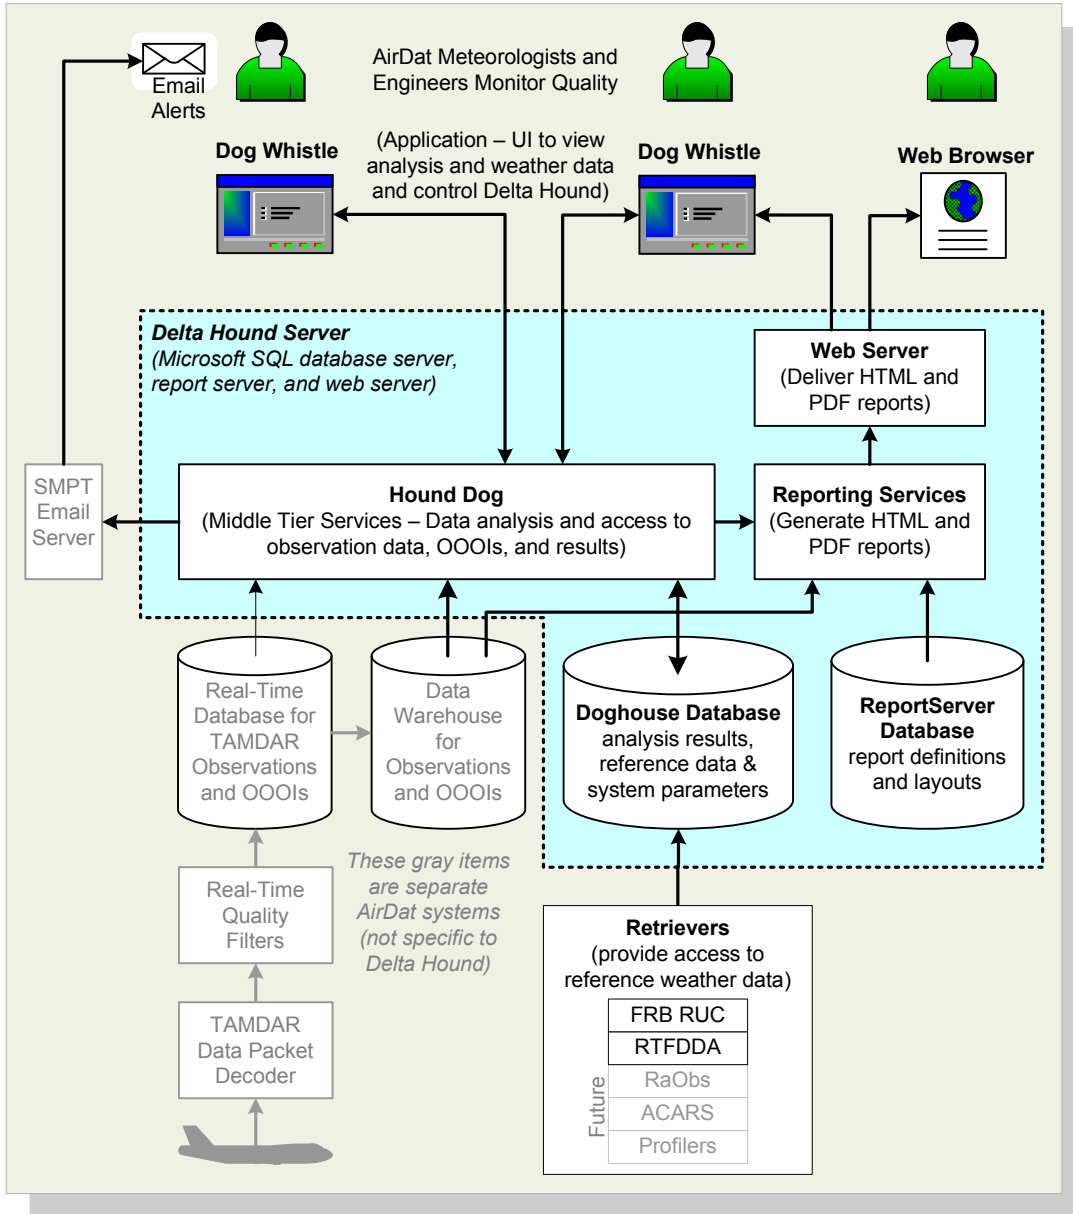


Figure A-16: Overview of the Delta Hound system architecture.