Volume 24 Number 9
September 2008

IFR®

The Magazine for the Accomplished Pilot

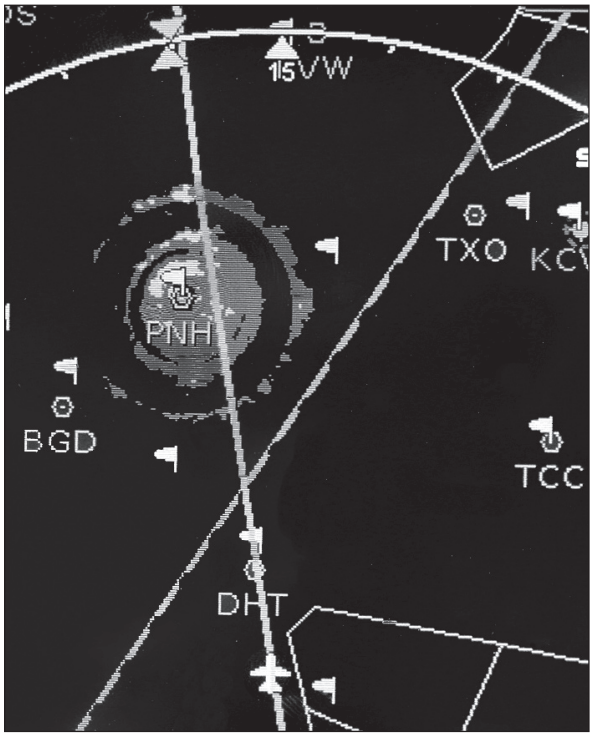

Are you really going to fly through that? ... page 9

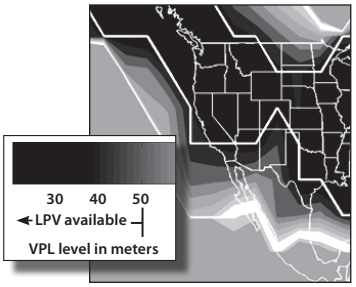

No LPV 4 U ... page 18

TRAIGHT-IN LANDING RWY 11 LOC (GS out)			
With Suvvo		Without Suvvo	
With Local Altimeter Setting			
MDA(H) 660′ (527′)		MDA(H) 720′ (587′)	
With Boston Altimeter Setting			
MDA(H) 720′ (587′)		MDA(H) 780′ (647′)	
	RAIL out / ALS out		RAIL out / ALS out
RVR 50 or 1		RVR 50 or 1	
RVR 50 or 1	1½	RVR 50 or 1	1½
RVR 60 or 1¼	1¾	RVR 60 or 1¼	1¾

Many minimums ... page 6

ALSO INSIDE THIS ISSUE ...

misses this until deep into the approach. Quick: balance a penny on your nose and subtract 2.5 from 6.1. I didn't think so. Do this calculation ahead of time. Write the numbers on the plate if you want.

Check the small table with time and vertical-speed info for several groundspeeds to catch a steep glideslope or to get an idea of what vertical speed will result in a stabilized descent to the TDZ for a non-precision approach.

Moving slightly right on the approach plate, review the approach-light configuration you can expect to see upon breaking out. In our case, "We have a MALSR system, with a PAPI on the left side of the runway."

Next re-brief minimums, verifying the correct number from what can be a field of DAs and MDAs. In our case, there are 18 minimums listed, combining permutations of five different variables: glideslope in service or not, local or remote altimeter setting, straight-in or circling, ability to identify the SUVVO fix, and aircraft speed.

Finally, commit to memory the items required inside the FAF: the DA or MDA, the determiner for the MAP, and the first step of the missed. In CFI speak, that's "How low, how far, and which way?" From the FAF inbound, all attention should be on the instruments, with no more reference to the plate.

Easier Done Than Said

As is the case with many things in aviation, it takes much longer to explain this process than to complete it. With practice, you should have no trouble completing a thorough briefing of an approach in under a minute. The process becomes more streamlined and automatic with practice and frees grey matter for the act of flying the approach itself, which is what it's all about.

Neil Singer is a CFI and turbine mentor pilot in the Northeast.

WEAVING RADAR MOSAICS

The NEXRAD image on your MFD combines various parts reality and best guesses. It may also omit some seriously nasty stuff.

by Scott Dennstaedt

Radar images beamed to our cockpit from satellites have become nearly indispensable. That little satellite receiver practically gives the pilot super powers. But glossing over the finer points sets you up for an unpleasant surprise someday.

We'll talk about the XM-based WxWorx radar image here. The radar image from WSI has some differences, but produces a nearly equivalent image most of the time.

Customized Radar

The weather radar product that eventually ends up on your portable or panel-mounted display is derived from NEXRAD, but it's not the same as the image you see on your favorite web site nor is it the same as the image seen by the radar operator at the local weather-forecast office. XM Satellite Radio may broadcast the image, but it's actually a radar mosaic produced by WxWorx.

Why not just pipe the raw National Weather Service (NWS) image down to your receiver? Because the image would contain many non-precipitation returns referred to collectively as ground clutter, making it difficult to interpret. The pilot wants to see only real precipitation. The raw radar data is generated by a network of NWS radars, which means you'd be required to select specific radar sites as you flew your route.

A better solution is a single, seamless image, called a radar mosaic, for the whole country. WxWorx, based in Huntsville, Ala., collects the data from 142 NWS radar sites into one continuous radar mosaic. They also attempt to improve the raw NEXRAD data by filtering out undesired radar returns, which means the raw data

Below: *A tornado over Amarillo, perhasp? Sometimes ground clutter slips through the filters. The circular pattern shown here over the Panhandle VOR (PNH) is a perfect example of what you'll typically see. The skies were essentially clear as we flew over.*

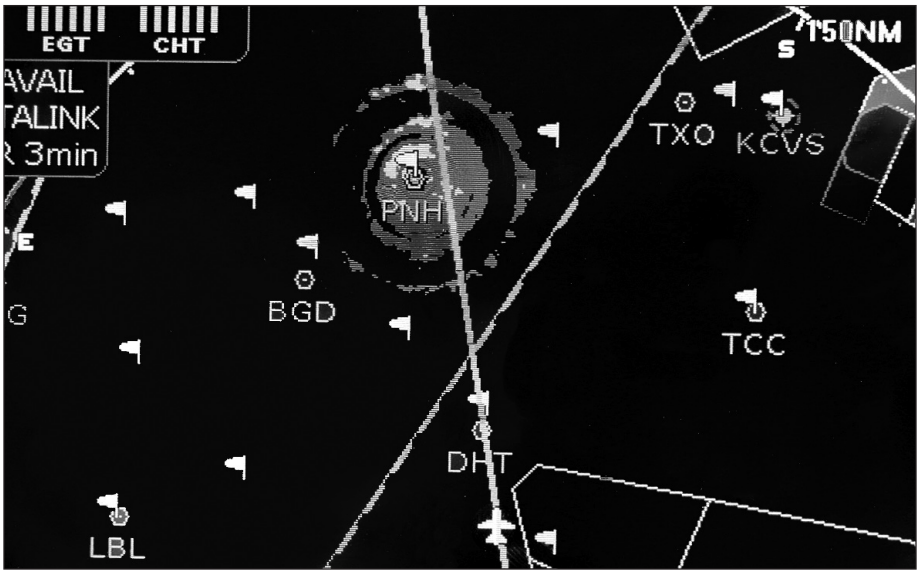

must go through a quality control and filtering process before they can be magically woven together.

Nexrad Backstory

In 1988, the NWS began to install the "next generation" radar that included a new feature called Doppler radar. Doppler radar allows forecasters to map the wind field in addition to precipitation. This Weather Surveillance Radar 1988 Doppler (WSR-88D) remains as our operational radar today.

Each of the 142 WSR-88D radars across the contiguous U.S. rotates continuously, sending out a beam of energy and listening for any of the energy to be reflected back to the antenna as it rotates. The returned energy is measured and the base data is stored.

The WSR-88D must make multiple revolutions or scans at different elevations or tilts. The number of elevation scans depends largely on the scanning strategy employed by the radar operator at the NWS local office. A complete set of scans at all elevations is called a volume coverage pattern (VCP).

Currently, the radar operator can choose one of four scanning strategies. The scanning strategy chosen depends on the kind of weather that's expected. Sensitive scanning strategies (called clear-air mode) are used when no precipitation is expected or when drizzle or light snow might be anticipated. There are two clear-air mode strategies and they use elevation angles from 0.5 degrees to 4.5 degrees. A complete, full-volume scan normally takes 10 minutes.

When heavier precipitation is expected, the radar operator will typically employ one of two precipitation-mode strategies. These strategies use more elevation scans (from 0.5 degrees to 19.5 degrees), but take less time to complete (five or six minutes for the full-volume scan).

Depending on the scanning strategy, each volume scan consists of a specific number of elevation scans. Each elevation scan provides a product called base reflectivity data. Note that "base" does not mean "lowest." Base reflectivity is available for every elevation angle. (It's called base reflectivity data to distinguish it from other base-data derived products.) It is the 0.5-degree base reflectivity data—the lowest elevation scan—that you normally see on the NWS web site and therein lies some confusion.

In addition to the 0.5-degree base reflectivity, the NWS produces a composite reflectivity product, which can also be viewed on your favorite web site. Composite reflectivity depicts the maximum reflectivity found on any constant-elevation sweep of a volume scan. That means adjacent areas of the composite image could contain reflectivity values projected from different heights within the storm.

Rain Only, Please

The data broadcast to your satellite receiver is also a composite reflectivity image, but it's one that WxWorx

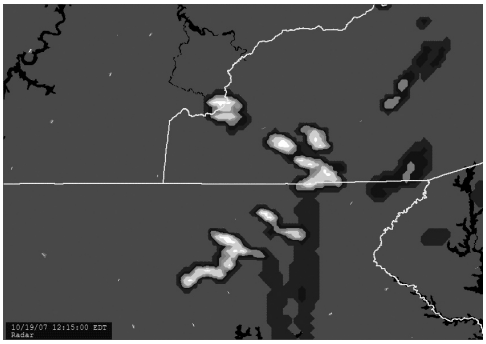

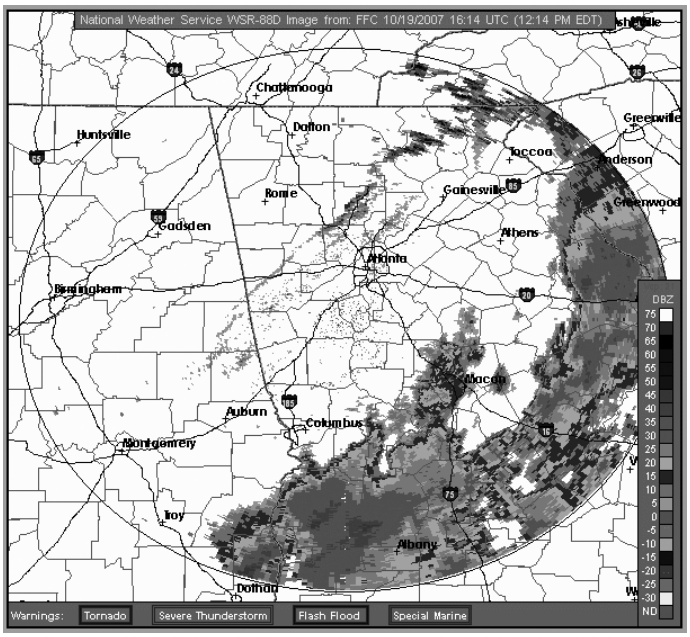

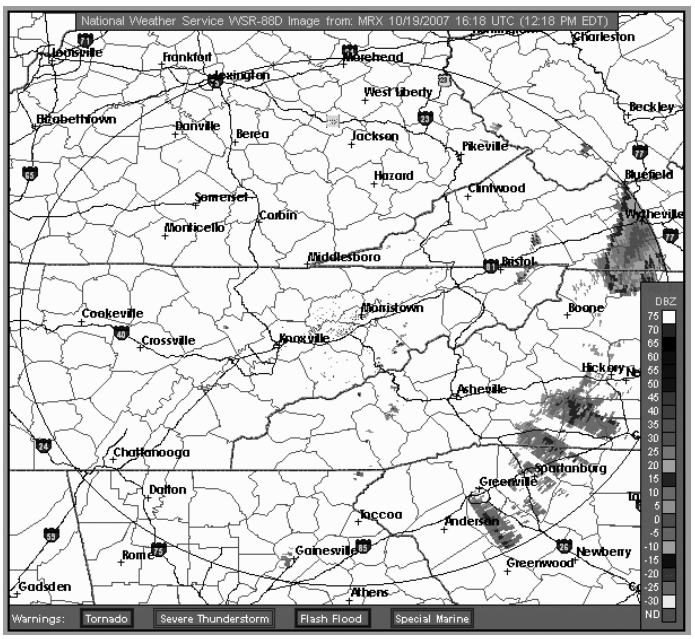

Anomalous *propagation (AP) can sneak through the WxWorx filters showing what appears to be intense convective cells (**left**). The Atlanta NWS radar (**lower left**) was the source of the AP in western North Carolina (just north of Toccoa), but a quick check of the Knoxville NWS radar (**lower right**) showed no returns. The precipitation was bogus.*

builds from the base data from each elevation scan. Some of the WxWorx process is "secret sauce" that they won't describe in detail. We know roughly what happens, however.

WxWorx takes the raw data from the NWS and places it through a series of filters in an attempt to remove non-precipitation radar returns—birds, insects, bats, dust and debris, military chaff, anomalous propagation, etc. This process isn't perfect. Real precipitation can be filtered and junk returns can sneak through.

The first filter is a new, manual step introduced earlier this year that employs a gross mask. Forecasters at the facility in Huntsville create fences around areas highly unlikely to contain real precipitation. Any base data that appears in this area gets filtered out. This will remove all of the clutter in these regions, but it will also remove real precipitation in rare circumstances.

Over-Zealous Filtering

WxWorx employs additional automatic filters, and they take any returns under 15 dBZ and toss them.

Unfortunately, low-power returns (<15 dBZ) could indicate larger cloud drops or ice crystals in a thunderstorm anvil. This is just one place were filters can go bad.

Most issues with "bad" data in your datalink radar are from non-precipitation returns sneaking through the filters. This tends to be more frequent during the evening and early morning when clutter tends to be the highest. Light returns may appear in a circular pattern around the radar site. It may be mixed with real precipitation or occur in clear air.

Anomalous propagation (AP) normally gets filtered out, but this occasionally slips through. Unlike nuisance ground clutter, AP can produce returns that look remarkably like real thunderstorms. AP is normally a result of a strong temperature inversion near the surface super-refracting the radar beam, causing it to strike the ground 100 or more miles away from the radar site. The returns

(continued on page 22)

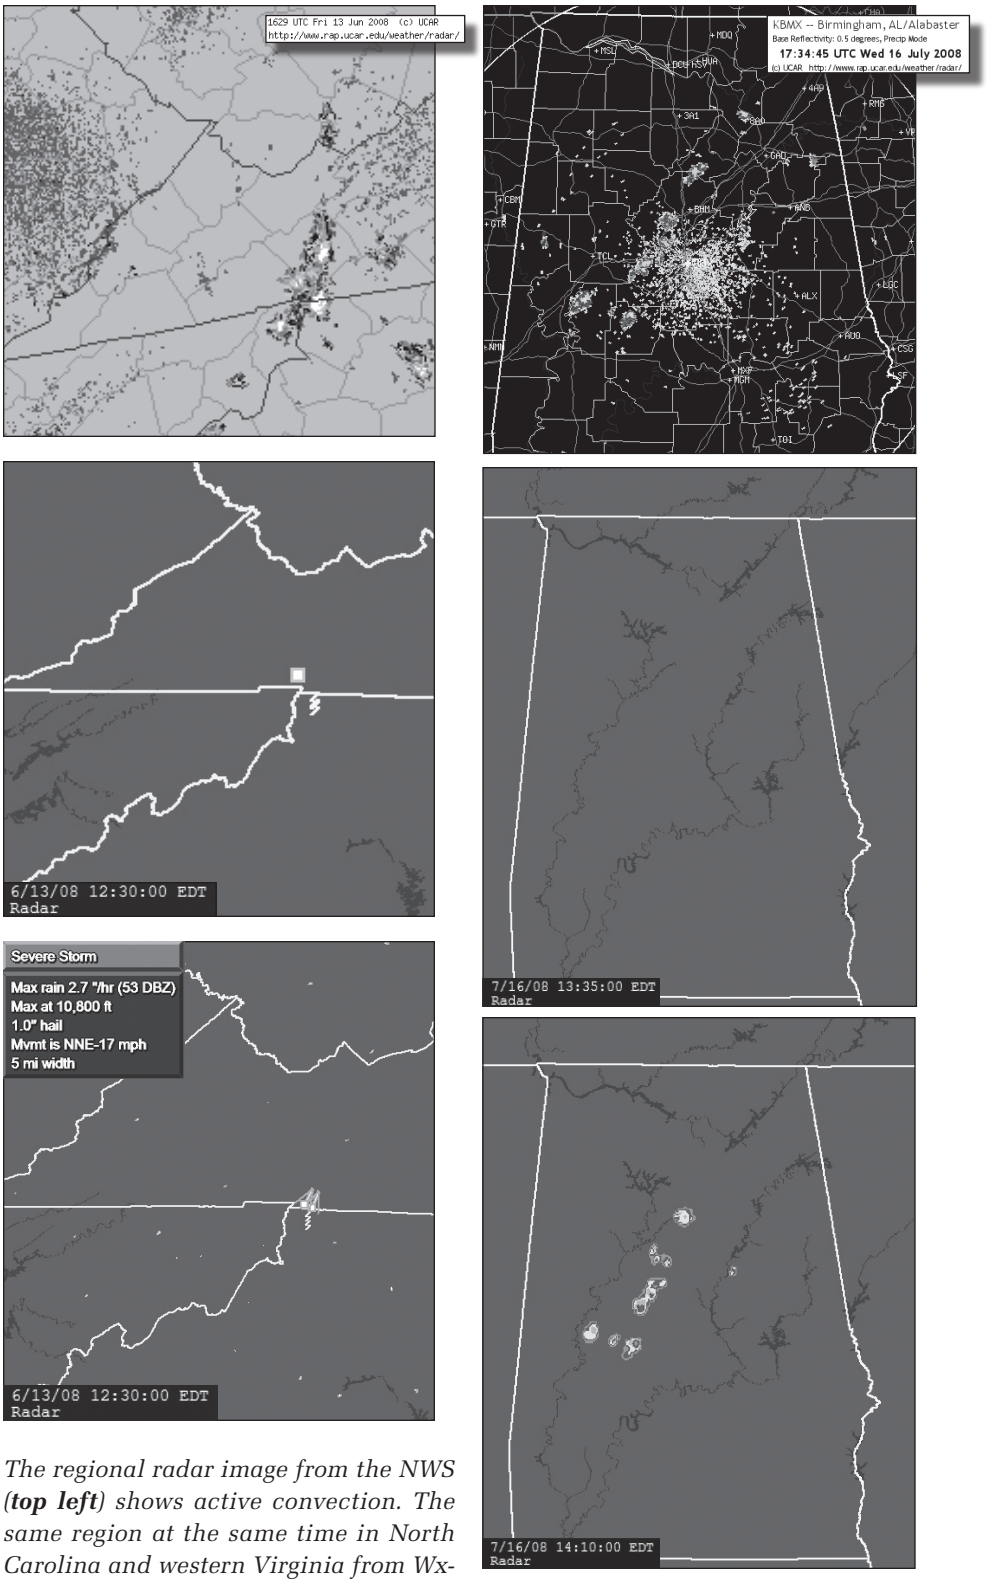

*The regional radar image from the NWS (**top left**) shows active convection. The same region at the same time in North Carolina and western Virginia from WxWorx on Wings (**middle and bottom left**) shows lightning and Storm Cell Information and Tracking (SCIT) markers, but no radar returns—even though the storm is really there. A worse case is shown with the radar from central Alabama (**top right**). The radar shows some clutter, but also real cells. The WxWorx image from the same time (**center right**) doesn't show anything. About 35 minutes later (**bottom right**), the cells just appeared, as if from nowhere. The only logical explanation is that WxWorx forecasters left the gross-mask filter on too long and filtered out the real precipitation.*

clog the frequency with useless information. With light traffic, a quick "O'Hare tower, Twin Cessna One Two Three Four Quebec is at 27 left, IFR, ready for takeoff," would be appropriate. But if the controller is up to his eyeballs in traffic, the "ready" call is a needless distraction.

Take a look at the traffic picture and try to guess what the controller's response to your call will be. If the answer is "roger," as it would be if you call "ready-in-sequence" with six airplanes ahead of you, hold your peace until it appears that you'll get a more positive response. The same applies when you're number one for the runway with traffic on final. Wait until the controller will say something other than "hold short."

Calling On Arrival

Arriving at an airport on an IFR flight plan does have its advantages; the Tower controller should know you're coming, and on what approach which runway. But a simple "Three Four Quebec's with ya" won't make him happy. He may be predicating IFR separation between you and departure on your position report, so give him what he needs: "Dupage tower, Twin Cessna One Two Three Four Quebec, over ARNIE on the ILS Runway 10, low approach, then VFR southwestbound."

If you'll be making a full-stop landing, the "who" and "where" will usually be sufficient, but if you're at an airport where practice approaches are common, adding the "what" (full stop, for instance) won't hurt. When your landing clearance includes instructions to hold short of another runway, acknowledge your intention to comply with these instructions. A laconic "roger" could come back to haunt both you and the controller if there's a misunderstanding.

If you don't feel comfortable landing in the distance available to the intersection, advise the controller immediately that you'll be unable to hold short. Phrases such as "we'll do our best" don't provide the assurance necessary to allow simultaneous intersecting runway operations.

On the ground and slowing, pay attention to the tower controller's instructions. If there's a runway to cross between you and your destination on the field, facility policy may require that the tower controller keep you on his frequency until you cross that runway. Wait for the tower to tell you to contact ground before switching.

One thing many pilots fail to realize is that once you've turned off the active runway, you have effectively ceased to exist as an IFR airplane. Up until now, there have been handoffs from one controller to the next, so each controller pretty much knew who you were, where you were, and what you wanted before you talked to him.

Not so with ground control at your arrival airport. To him or her, you're a brand new customer, maybe an IFR arrival, maybe a VFR arrival, or maybe just an airplane requesting taxi from one point on the airport to another. Maybe you've landed here every day for the last 20 years and taxied to the Whiz-Bang Aviation tiedown area, but maybe the ground controller just got here yesterday, so he won't know that. Give him the information that he needs.

Back in the Cab...

Shortly after TransOcean 462 cleared the runway, I heard another voice on the frequency:

"O'Hare ground, Cessna One Seven Zero Lima Bravo, clear of 22 Right at the outer taxiway, request taxi to Butler."

Looking out of the cab, I saw a pristine Cessna 170 at precisely that position, looking painfully out of place amongst the multitude of taxiing airliners. It didn't really matter what type the airplane was, though, I could tell by the way he handled the radio that I was talking to a pro.

This article first appeared in the February 1991 IFR *when Denny Cunningham was still a tower controller at O'Hare International in Chicago. Denny is now retired, but his suggestions are still on the mark, 16 years later.*

QUIZ ANSWERS *(questions on page 16)*

1. **d.** This is probably the best answer. The AIM suggests an optimal rate of climb in all situations until 1000 feet below your assigned altitude and then 500 fpm thereafter. You're supposed to tell ATC if you can't make 500 fpm. Just remember the controller is watching your rate as you start to see how you're doing and planning based on that. Don't start fast and then slack off.

2. **a.** Again, this is the best answer. It's possible that the facility has a waiver to allow faster traffic. In the end, complying with FAR 91.117 is your responsibility not ATC's. They may be asking you to speed to cover their own spacing issues. Not technically legal, but it's been known to happen in the real world.

3. **True.** As long as no other restriction was included in the clearance, when and how you change altitudes to meet the crossing altitude is up to you.

WEAVING RADAR MOSAICS

continued from page 11

follow the same path back to the radar and are sometimes interpreted as high reflectivity values. AP can also be caused by a cold pool of air after the passage of thunderstorms.

The high anvil clouds on thunderstorms can produce returns that are not necessarily precipitation. After the thunderstorm grows to maturity, it spreads laterally at the mother of all inversions called the tropopause, creating the anvil shape. The NWS radar can see the ice crystals in these high clouds and may

produce low reflectivity returns that WxWorx does not typically filter out, since they are strong enough to pass through the low-dBZ filter. If you're flying under one of these anvils in clear air, the radar image might show you flying through light precipitation even though no precipitation is falling. This can make the thunderstorm appear bigger on your screen than it really is and is common around dissipating cells.

Gust fronts and thunderstorm-outflow boundaries are not precipitation echoes, but are clearly seen on the NWS radar image as low-reflectivity returns. They are filtered out by WxWorx since they are generally below 15 dBZs and occur in clear air most of the time. Gust fronts can produce significant straight-line winds and wind shear near the surface and can extend five or more miles from the main precipitation core.

While the gross filter is applied in locations that are *nearly* impossible to contain precipitation, on a few occasions unexpected precipitation has developed there. This was related to a latency between the time the filter was drawn, and the time it took to get into the WxWorx data stream. WxWorx has recently identified the cause of this, and claim they have adjusted their processes to overcome the latency. This is a serious issue so we'll continue to monitor it.

Old News

The radar data is broadcast to your receiver on a strict schedule. You see this time on your cockpit display. The important fact here is that the time you see isn't the age of the NEXRAD image. It's the time since the scheduled broadcast.

There are three potential areas of delay, or latency: the volume scan of the individual radars, the time it takes to composite and ship the mosaic to XM and the time it takes to uplink to the satellite. The last two are typically no longer than 30 seconds combined, depending on how much activity there is.

WxWorx broadcasts their composite image every five minutes. Their schedule doesn't match the NWS schedule of five, six or 10 minutes, however, so the WxWorx image may be built from two different NWS volume scans from some unknown mix of the 142 individual stations. In other words, if a scanning strategy consists of 14 elevations for a complete volume scan, WxWorx will use base data from the current volume scan (elevations 4, 3, 2, and 1) as well as the base data from the previous volume scan (elevations 14 through 5) to produce their own composite image. There's no way to tell what radar or elevation scan a particular "pixel" came from.

I believe that pilots should assume the image on their display right after it is refreshed could be four to eight minutes behind reality (assuming no broadcasts are missed). A cell that's clipping along at a significant rate could move four or five nautical miles by the time you reach the spot where the cell is shown. In a fast-moving system, what you see on your display may not match what you see (or can't see) out of your window.

In reality, WxWorx doesn't send an image, either. They broadcast data in a specific format. It's up to the vendor (such as Garmin) to build the final image you see. Garmin is free to display intensity in six shades of blue, as four different colors (e.g., green, yellow, red and magenta) or any method they desire. Garmin and other vendors simply document the details in their user manuals.

WxWorx produces an excellent product most of the time. Don't trust it with your life, though. Use it as secondary guidance and be prudently skeptical of what you see on the digital display. Your primary guidance still should be what you see on the other glass panel in your cockpit—the windscreen.

Scott Dennstaedt is a former NWS meteorologist and active CFI. His web site is http://avwxworkshops.com.