Table of Contents

Topic: Flash Floods

Click to jump to lesson

Hi, my name is Jill Hardy and welcome to this lesson on flash flood meteorology.

We have a guest speaker for this lesson: Steve Martinaitis of OU CIMMS at NSSL. But if you have any questions regarding the material, please feel free to contact me, or the RAC team. Our contact information will be on the next slide.

Back to Top

Lesson Objectives

- Identify the mesoscale and storm-scale variables that contribute to the flash flood potential
  - Precipitation Rate/Efficiency
  - Precipitation Duration
- Identify heavy rainfall using WSR-88D and Dual-Polarization technology

There are two main objectives with this lesson. The first is to identify the variables related to precipitation rate and duration that contribute to the flash flood potential at a meso-scale and storm-scale levels. The second objective is to identify rainfall signatures using the WSR-88D and the new dual-polarization technology.

Back to Top

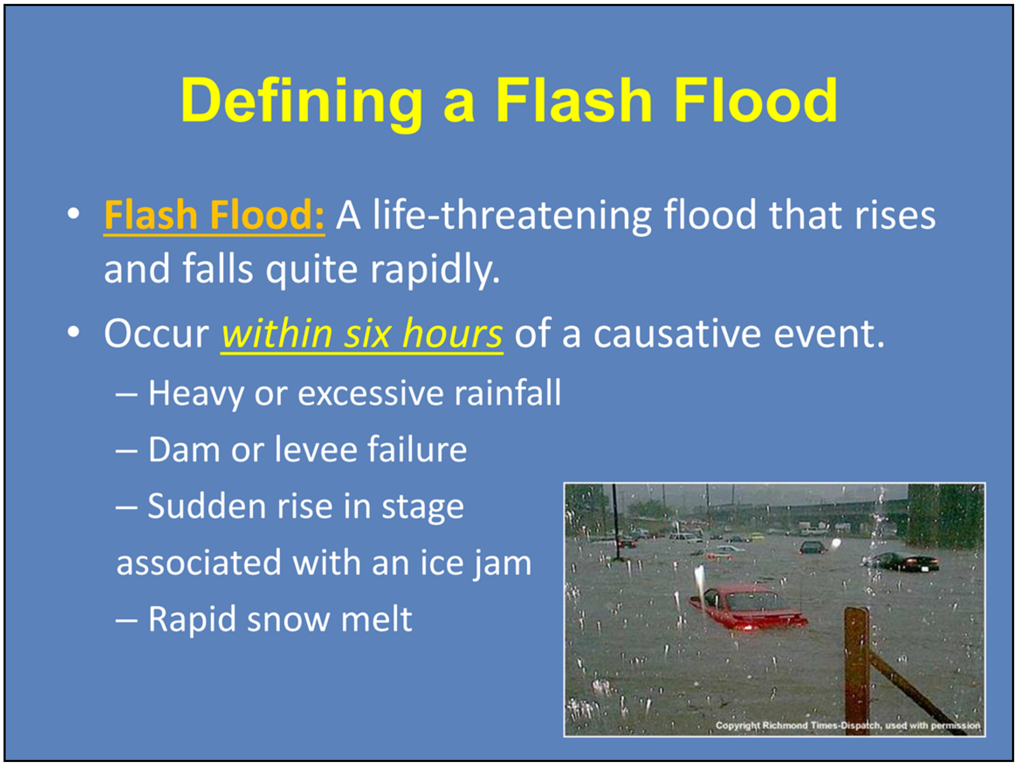

To guarantee that we are all on the same page, I want to make sure we understand how a flash flood is defined. Basically, it is a life-threatening flood that occurs quite quickly, i.e., within a six hour period. Flash floods can occur from a variety of events, such as heavy rainfall, dam failures, ice jams, or rapid snow melt.

For the purpose of these lessons, we will focus on flash flood events related to heavy rainfall.

Back to Top

Meteorological Ingredients

Precipitation Rate

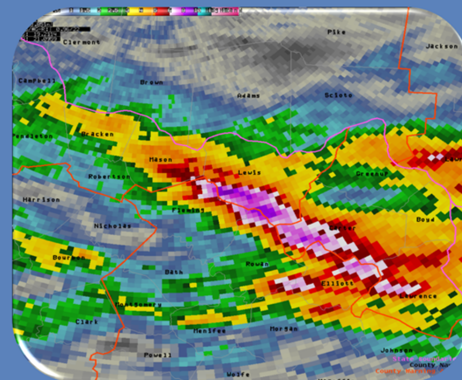

Precipitation Duration

When it comes to the meteorological aspects of flash flooding, the two most important things to consider are the precipitation rate and the precipitation duration. Let’s focus on the factors that influence the rate first.

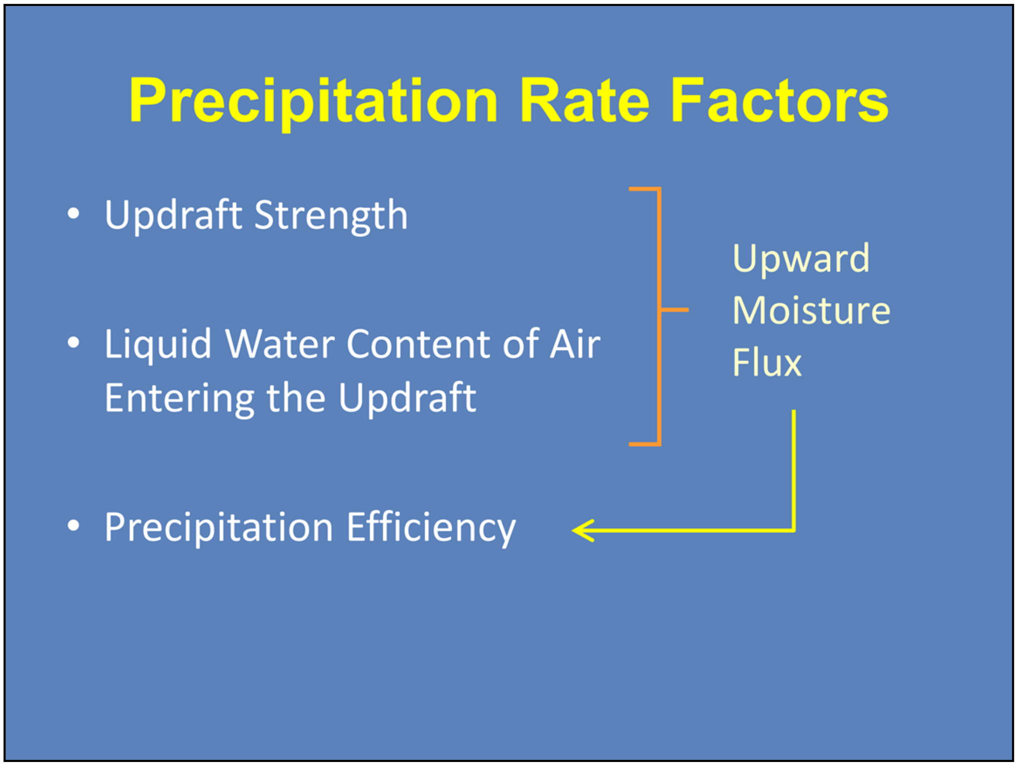

There are several factors that help determine the precipitation rate. The updraft strength and the liquid water content of the air that is entering the updraft contribute to the upward moisture flux into a storm. The percentage of that moisture flux that returns to earth as precipitation characterizes the precipitation efficiency of the storm.

Back to Top

Precipitation Efficiency

- Definition: Fraction of total moisture ingested by the updraft that falls back to the ground
- Dependent upon…
  - Updraft Strength
  - Moisture Profile
  - Warm Cloud Layer
  - Cloud Seeding

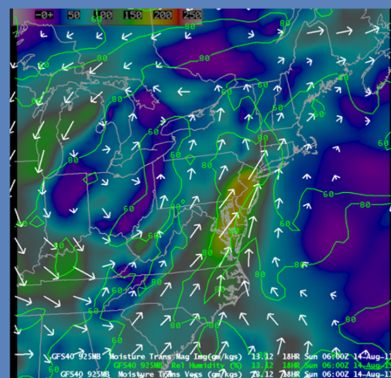

As you see here on the slide, precipitation efficiency is defined as the fraction of total moisture ingested by the updraft that returns as precipitation. Precipitation efficiency cannot be quantified in real time, so you will need to examine a number of factors to infer an efficiency. These factors include the updraft strength, the vertical moisture profile of the atmosphere, the depth of the warm cloud layer, and cloud seeding.

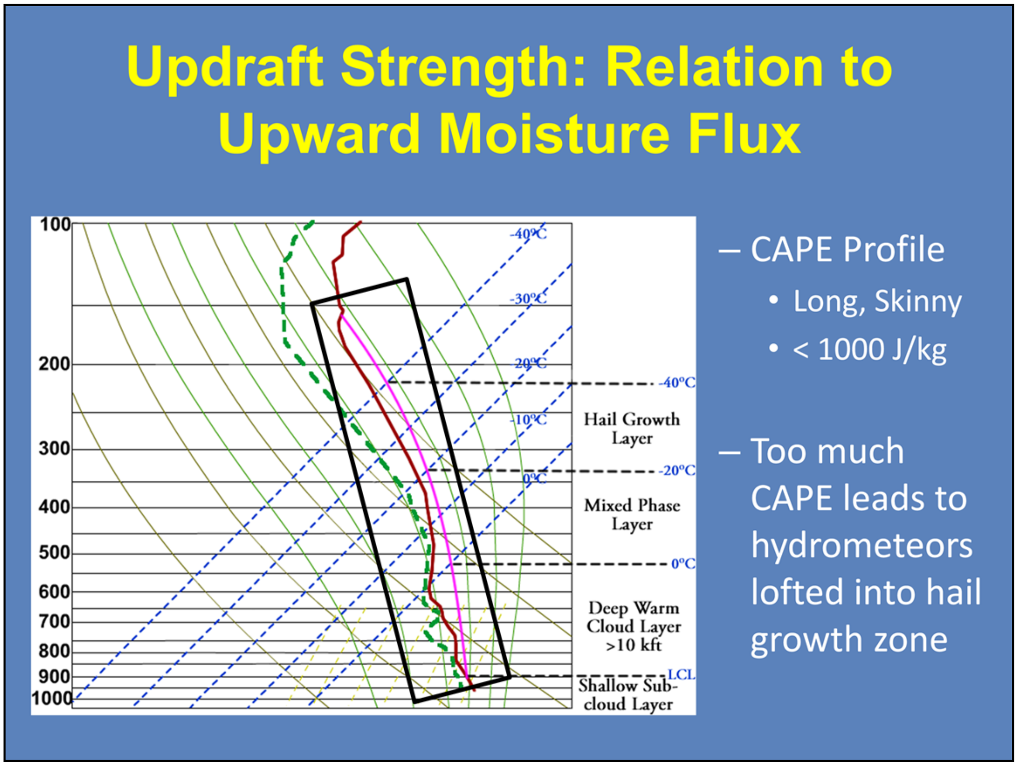

Let's start with the strength of the updraft. Shown here is the average atmospheric profile of flood and flash flood events. When considering the updraft strength of a convective storm, you would want to see a long and skinny CAPE profile. The amount of CAPE in the atmosphere should be under 1000 J/kg. Larger CAPE values will loft the hydrometeors ingested by the updraft into the hail growth zone.

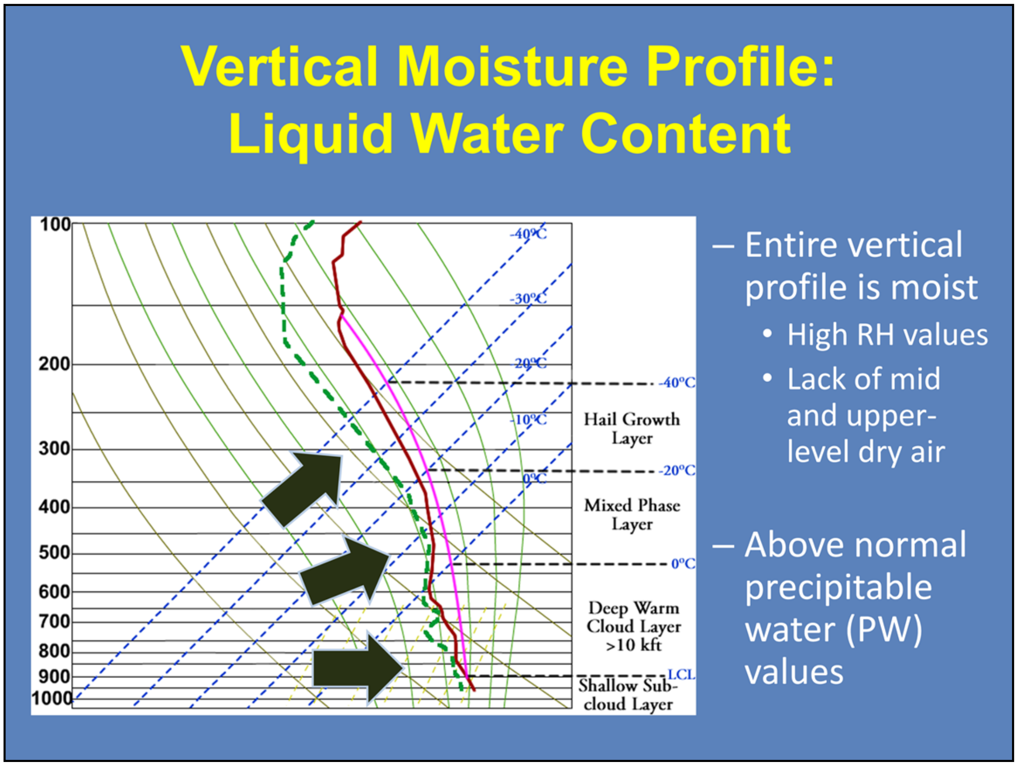

Now looking at the temperature and dew point profile of this sounding, you can see that it is very moist at all levels. Notice how there is a lack of dry air at the mid and upper levels. This is important when you consider the depth of water within a column of the atmosphere if all the water were precipitated as rain, otherwise known as the precipitable water (PW) value. Seeing above normal PW values is a good indication of how moist the atmosphere is.

So how do you determine what is an above normal PW value?

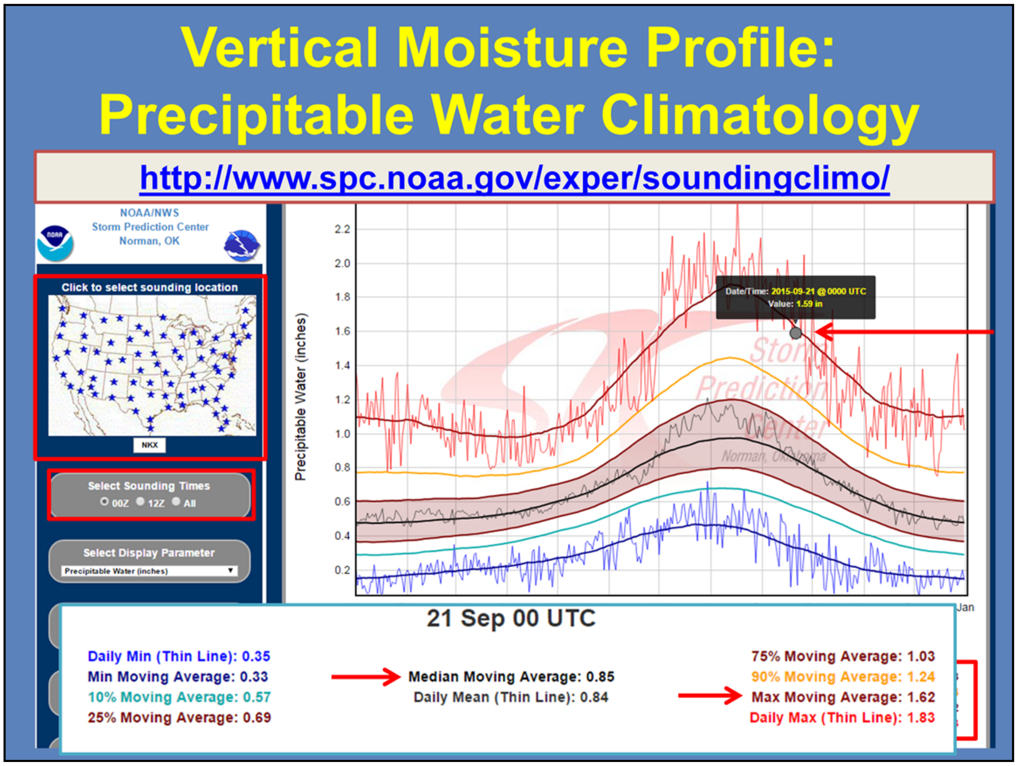

I'm going to briefly hop in here, as there's been a change since this lesson was created. Many of you may have been familiar with Matt Bunkers' Precipitable Water Climatology page. However, the SPC now hosts the point sounding climatologies, similar to the previous website.

Using the SPC site, comparing model or observed precipitable water (PW) values is quite easy. Begin by navigating to the desired sounding location, and select a sounding time. The plot now shows the daily minimum, several moving percentile averages, as well as the daily maximum for each day of the year.

Let's use this 00Z sounding climatology plot for the KNKX radar near San Diego. When I overlay the latest sounding information, we see the current value is 1.59 inches. This is near the maximum moving average of 1.62 inches for this day. Historically on this day, the median PW is 0.85 inches, so we are quite a bit higher than that.

Heavy precipitation events that lead to flooding and flash flooding have values that are above the 75th percentile and usually approach the 90th or maximum moving averages. In fact, for this example, the San Diego WFO had a flash flood watch in effect for the majority of their CWA. Use the URL to access the PW climatology page.

Alright, back to Steve!

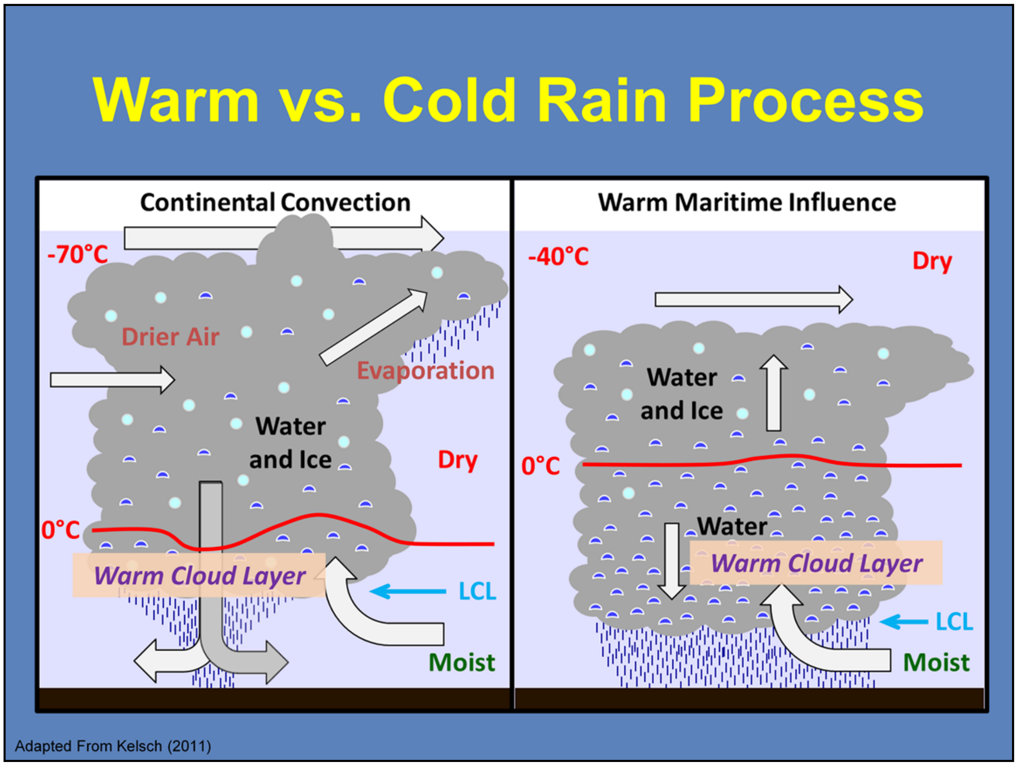

Now that we have analyzed the CAPE and moisture profile of the atmosphere, we can see how it helps determine whether a warm rain or a cold rain process is the predominant precipitation production method. Recall that precipitation forms through collisions and coalescence within a warm rain process while deposition and the Bergeron Process (the collision of ice crystals) define a cold rain process.

Looking at convection derived from a continental airmass, you can see that the LCL is relatively high while the in-cloud freezing level is quite shallow. The vertical separation between the LCL and freezing level is defined as the warm cloud layer. This is where warm rain processes occur. However, the warm cloud layer is generally not very deep with this type of convection. Within a strong CAPE environment, hydrometeors will be lofted beyond the warm cloud layer, where they will become frozen (resulting in the formation of hail) and become subjected to evaporation due to mid and upper level dry air entrainment. This region is where the majority of the hydrometeors undergo cold rain processes.

Now focusing on the convection influenced by a warm maritime airmass, you notice that the LCLs are relatively low, and the in-cloud freezing level is much higher. Therefore, you have a greater warm cloud layer. The weak CAPE profile allows for the majority of the hydrometeors to remain below the freezing level. The moist vertical profile also helps in diminishing the effects of evaporation and dry air entrainment. Here, warm rain processes will dominate precipitation production.

Back to Top

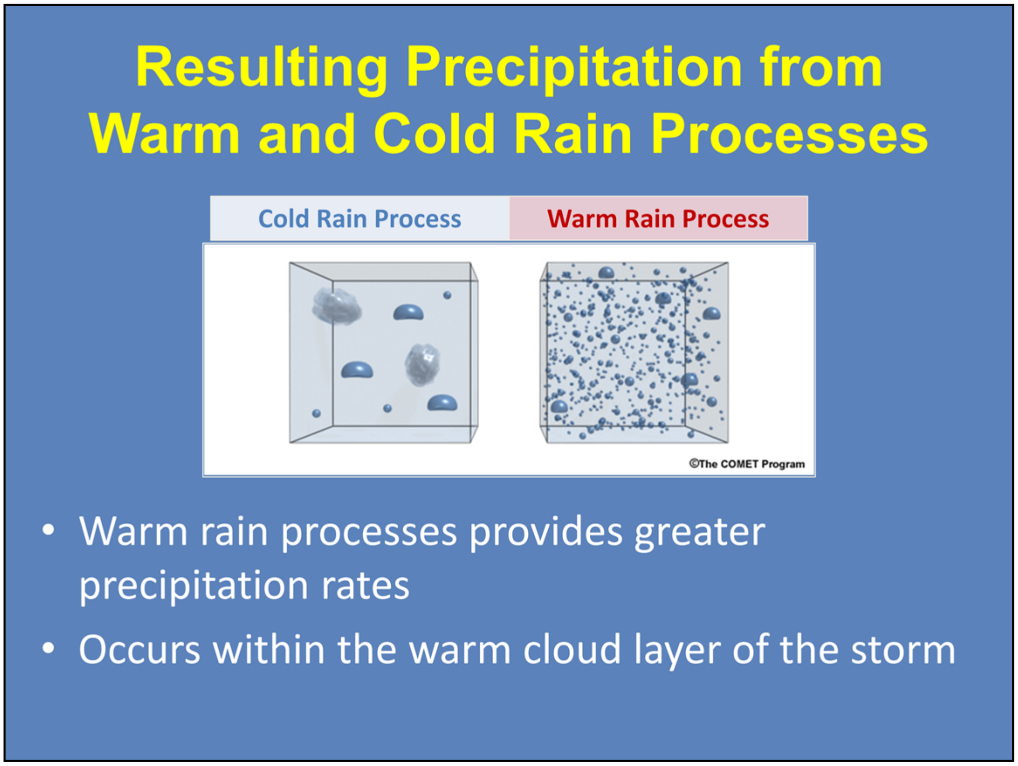

Comparing the resulting precipitation at the surface, you can see the dominant cold rain processes from the continental airmass yields a small quantity of rain drops that are generally large in size and can also include hail stones. The dominant warm rain processes in the maritime airmass has a substantial quantity of raindrops. So, the warm rain process results in a greater precipitation efficiency and greater precipitation rates.

In the example on the previous slide, you saw how the CAPE and moisture profiles influence the amount of hydrometeors that reside in the warm cloud layer, and thus, could undergo warm rain processes. Which leads to the next set of questions… How do we calculate the warm cloud layer? And how deep of a warm cloud layer do you need for precipitation rates that could potentially yield flooding?

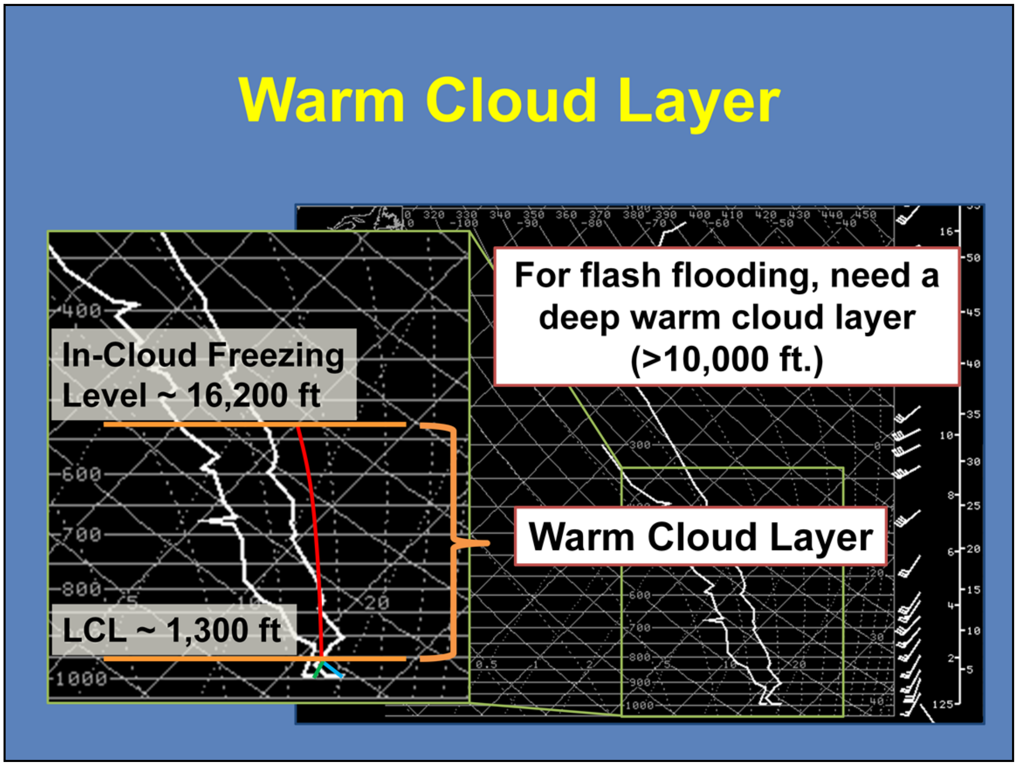

To calculate the warm cloud layer, we will use this sounding from Grey, ME during a summer-time convective event that had some flash flooding. Focusing on the lower and mid levels of the atmosphere, you would start by finding the LCL. Follow the dry adiabatic lapse rate from the surface temperature and the saturation mixing ratio from the surface dewpoint until the two lines meet. The LCL for this case is about 1,300 feet.

From the LCL, follow the moist adiabat up to the freezing level. We choose the moist adiabat because that should be "in-cloud" and also where the warm rain process (collisions and coalesence) is occurring. In this case, the freezing level is around 16,200 feet. The difference in height between the LCL and in-cloud freezing level will be our warm cloud layer. Having a *deep* warm cloud layer is very important for flash flood forecasting. A warm cloud layer over 10,000 feet is considered deep.

For this example, our warm cloud layer is approximately 14,900 feet.

Cloud Seeding

- Definition: Jump start of precipitation product via the ingesting of hydrometeor embryos
- Inter-Storm Seeding
  - Increases upward Moisture flux
  - Increases environmental humidity

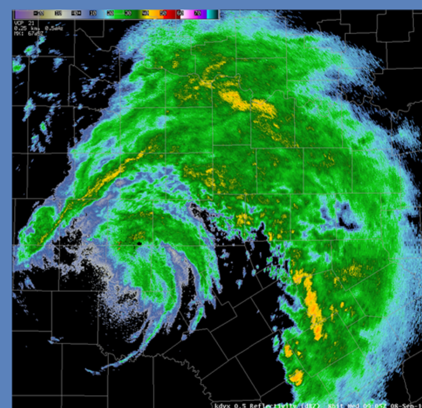

Another process that increases precipitation production is cloud seeding. We will focus on inter-cloud seeding here. This is the process where precipitation production is jump started by the updrafts ingesting hydrometeors from other storms. This will help increase the upward moisture flux and increase the local environmental humidity. In this example, an intense rain band forms with the remnants of Tropical Storm Hermine over central Texas. The combination of the tropical environment and the inter-cloud seeding enhanced rainfall production in an already efficient precipitation environment. Widespread rainfall totals of 6-10 inches were common with system.

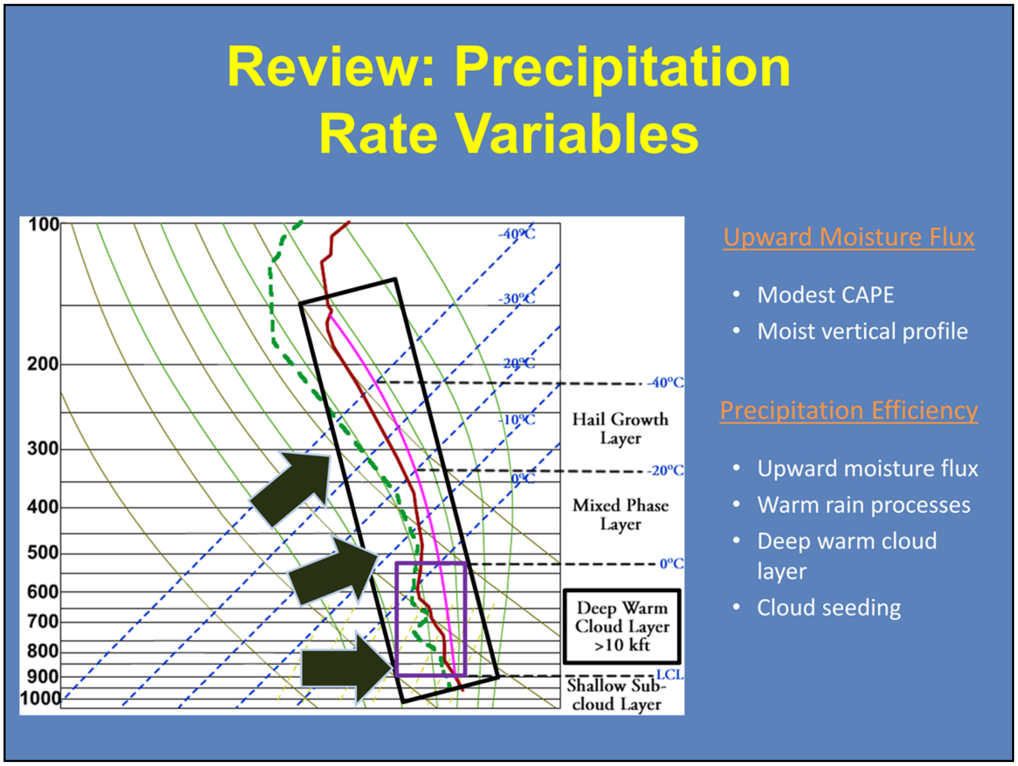

Let's quickly recap what influences the precipitation rate. First, you have to consider the variables that go into the upward moisture flux of a convective storm, such as a modest CAPE profile, generally under 1000 J/kg, and a moist vertical atmospheric profile.

The fraction of the upward moisture flux that becomes precipitation defines the precipitation efficiency of the storm. Along with the upward moisture flux, recall that warm rain processes provide greater precipitation rates. Warm rain processes occur within the warm cloud layer. Remember that for a greater flash flood potential, you would like to a deep warm cloud layer of 10,000 feet or greater. You also have to consider inter-cloud seeding to increase precipitation production.

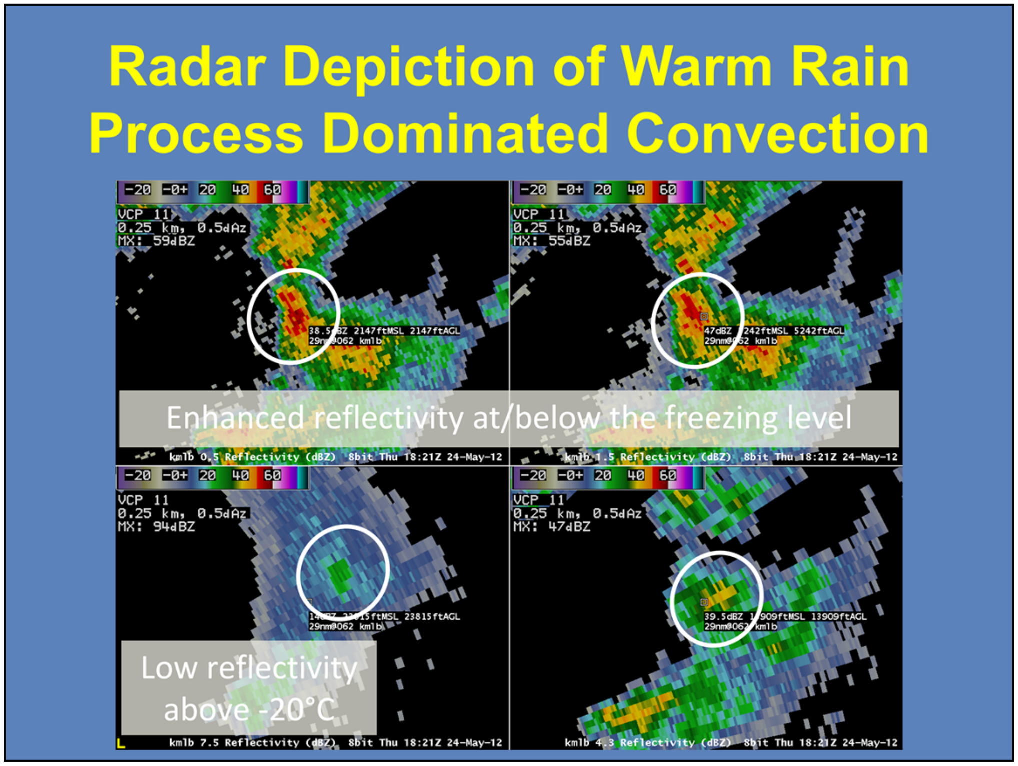

So what does convection dominated by the warm rain process look like on radar? Honestly, not much unless you know what you are looking for. The two main characteristics of convection dominated by the warm rain process are enhanced reflectivity values at or below the freezing level and low to non-existent reflectivity values above the -20°C level.

Using this example from the Melbourne, FL office, the top two images show reflectivity values between 50-60 dBZ below the freezing level. In the bottom-right panel, the 4.3° tilt scans the storm at 13,900 feet, just a few hundred feet below the freezing level. Here, there are very few pixels that meet or exceed 40 dBZ. The bottom-left panel shows the storm at the 7.5° tilt near the -20°C level. Reflectivity values here are below 25 dBZ. The storm does not exist on higher tilts.

This series of images shows what is called a low-echo centroid signature. This is where the majority of the precipitation core lies below the freezing level. The combination of this type of radar signature and a moist, slightly unstable environment should clue you in to warm rain processes being dominant here.

Back to Top

Identifying Heavy Rainfall using Dual-Polarization

Characteristics

- 50 dBZ < Z < 60 dBZ
  - 40 < Z < 55 dBZ for tropical environments
- 2.0 dB < ZDR < 5.0 dB
  - 0.5 < ZDR < 3.0 dB for tropical environments
- CC > 0.96
- KDP > 1.0 deg/km

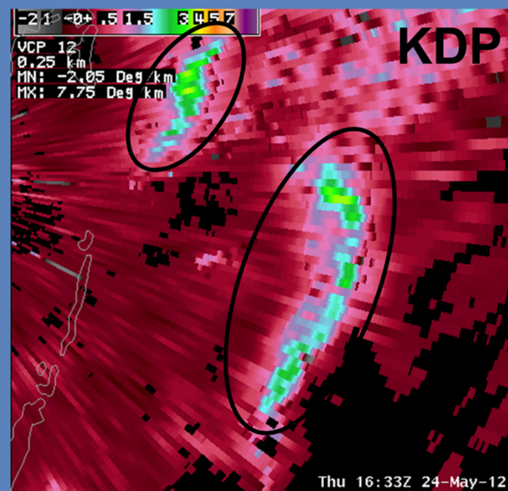

With the addition of dual-polarization technology, the new algorithms can help pinpoint areas of greater precipitation rates. This signature is from the Miami, FL radar and was related to a tropical disturbance that eventually became Tropical Storm Beryl.

Starting with the reflectivity (Z), you would look for areas of enhanced values, generally in the 50-60 dBZ range; 40-55 dBZ for tropical environments. Here, we are highlighting two areas of enhanced values.

Now examining the differential reflectivity (ZDR), the difference between the horizontal and vertical reflectivity factors, you would look for ZDR values between 2.0 and 5.0 dB, and 0.5-3.0 dB for tropical environments. Remember, there is a strong relationship between the raindrop size and ZDR where the greater the ZDR values, the larger the raindrop diameter. Since we are dealing with a tropical environment in this case, the ZDR values are around 1.5 dB. Combine that with the high reflectivity values, you have a lot of small rain drops here.

Moving on to the correlation coefficient (CC), you should see very high values (above 0.96). This means that the type of precipitation that is being sampled is uniform. As you can see here, the areas that had the greater reflectivity have a CC of around 0.99, meaning all the precipitation here is rain.

Finally, values of the specific differential phase (KDP) should be above 1.0 deg/km. Higher KDP values can mean larger rain drops or a larger concentration of rain drops. Since we know this is a tropical environment and the ZDR values suggest small rain drops, then this means we are dealing with a larger concentration of rain drops, and thus, greater precipitation rates.

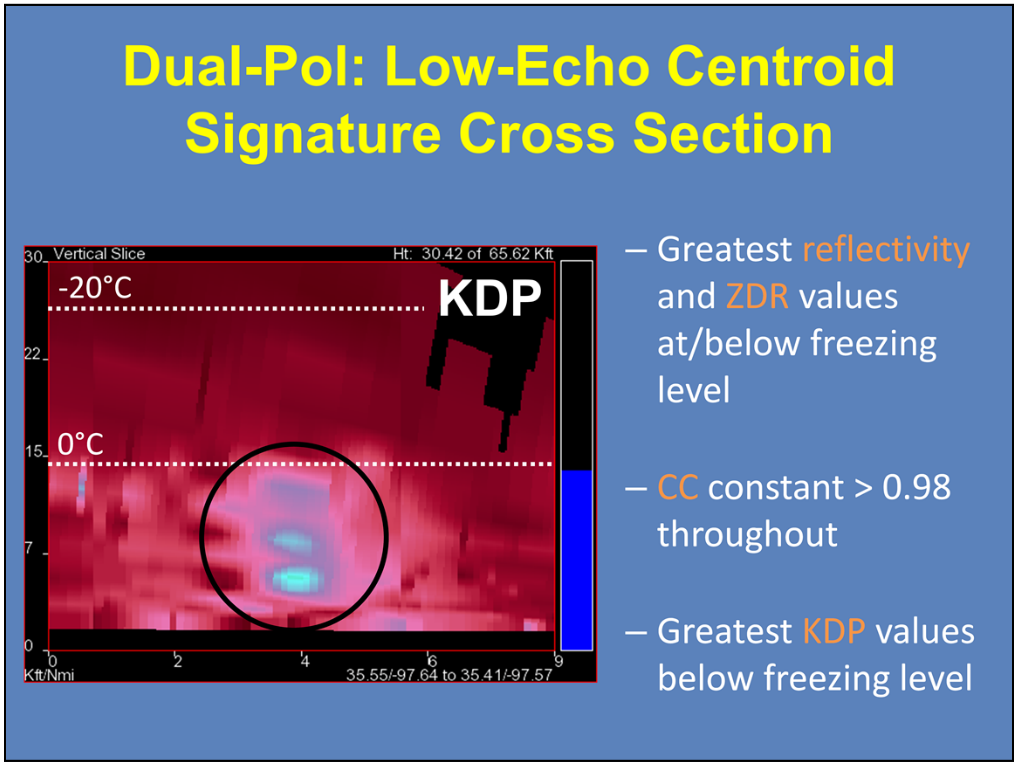

Now that you have seen what a warm rain process dominated storm looks like with base reflectivity and with the dual-pol products, we will now take a cross-section through a low-echo centroid signature. This example will look at a specific storm that was part of a system that produced significant flash flooding on the north side of Oklahoma City, OK.

As you saw in the four panel image earlier in the presentation, most of the enhanced reflectivity values lie at or below the freezing level and low reflectivity values exist near and above the -20°C level. The greater ZDR values, which represent rain drops, also lie below the freezing level. The very low ZDR values above the freezing level can represent very small water droplets, ice crystals, and/or hail.

The CC values are constant throughout the vertical profiles with them ranging from 0.98 to 0.995. The values closer to 0.98 (the darker purple shading with a slight orange tint) represent all rain with slightly larger drop sizes. Finally, the greater KDP values exist below the freezing level, showing where the greatest concentration of rain drops are occurring.

Back to Top

Dual-Pol: Identifying Heavy Rainfall with Supercells

Characteristics

- Z > 55 dBZ
  - Hail/Rain mixture
- ZDR can be anything
- CC < 0.96
- KDP > 1.0 deg/km
  - Normally most extreme
  - KDP not shown when CC < 0.90

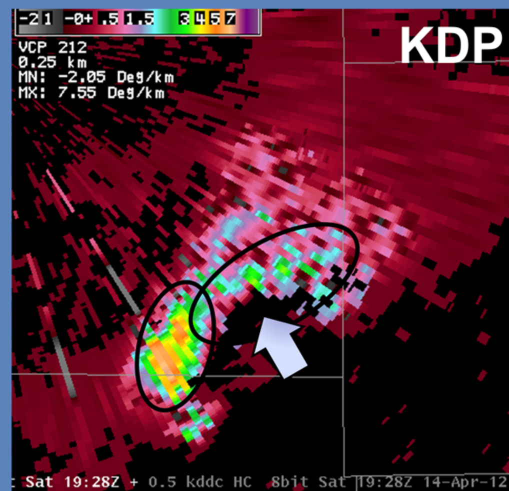

Since we have looked at what an efficient rainfall producer would look like with radar, let's take a look at what an inefficient storm would look like in dual-pol. For this example, we will use a supercell viewed from the Dodge City, KS office during the April 14, 2012 outbreak. Supercells can produce heavy rainfall, but you would need to examine the characteristics and motion of the storm to determine its flash flood potential.

Starting again with reflectivity (Z), you would look for areas that are greater than 55 dBZ. Here, we have highlighted two separate areas within this supercell. These areas of enhanced values are probably areas of hail/rain mixture.

Now starting with the dual-pol products and differential reflectivity (ZDR), it should be noted that the ZDR values can be anything because of hail contamination. Severe hail can bring ZDR values to near 0 dB while water coated hail can have values up to 6 dB.

Since we are dealing with non-uniform precipitation types, correlation coefficient values will be below 0.96 in areas of rain/hail mix. Here, we see values ranging between 0.9 and 0.95, with some lower values within the forward flank downdraft. Now overlaying the hydrometeor classification algorithm (HCA), you can see where the radar is seeing the hail/rain mixtures in red.

Finally, looking at the specific differential phase, you would see values greater than 1.0 deg/km here, and you do in both of the highlighted areas. Some of the more extreme values, like the area of 4.0-7.0 deg/km near the rear flank downdraft, are where the greatest rainfall rates are occurring, but some values could be a result of water coated hail. It is important to note that KDP values will not display in areas of CC less than 0.90.

Back to Top

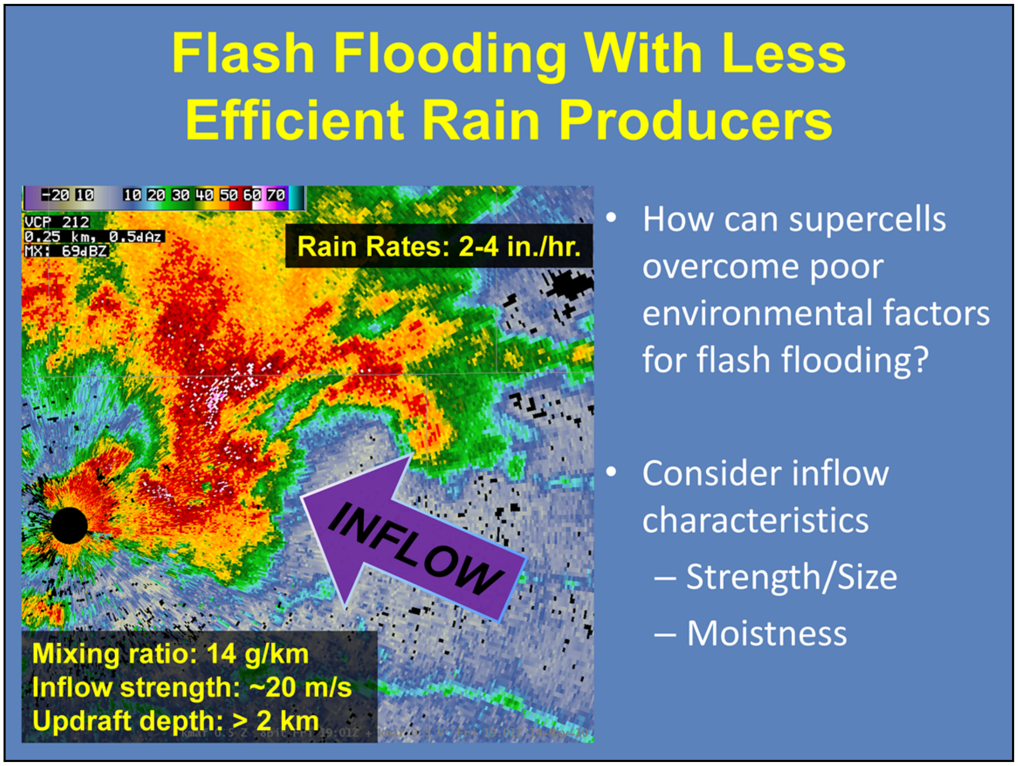

The key to high precipitation rates with supercells is understanding why supercells are such a threat despite poor precipitation efficiencies and at times swift movement. You just saw the dual-polarization characteristics of a supercell. However, there are some environmental factors to consider. In this case, we will focus on how much air it is ingesting and how moist is the inflow. Looking at this high-precipitation (HP) supercell near Midland, TX, this storm is in an environment where the profile is moist up to 700-mb and is considerably dry above it (The warm cloud layer is about 9,000 ft.). So, this storm is undergoing processes such as dry air entrainment and evaporation.

However, the storm has a very strong, moist updraft with it. It is ingesting very moist air (mixing ratio of 14 g/kg) at about 20 m/s. The updraft is wide and has a vertical depth of over 2 km. So, even though this supercell existed in an environment characterized by dry air above 700 mb, the storm produced rain rates of 2-4 inches per hour and fatal flash floods in the city of Midland. This shows that the factors that give HP and classic supercells a lower precipitation efficiency can be balanced by large values of moisture inflow and why supercells can produce high rain rates and flash flooding.

Precipitation Duration Factors

- Residence time of precipitation over a location
  - Rainfall Area
  - Storm Motion
  - Boundaries
  - Training Storms

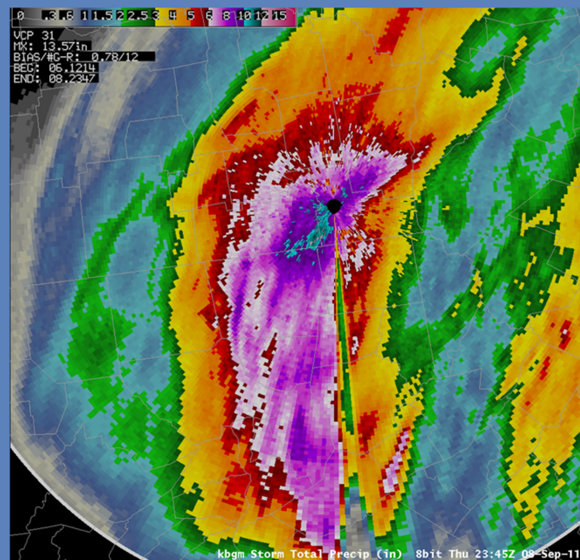

Now that we have examined the meteorological variables that influence the precipitation rate, let's examine the other meteorological factor that can influence the flash flood potential: precipitation duration.

When we talk about duration, we are talking about the residence time of precipitation over a location. There are a number of things that affect the duration of rain over a specific area, such as rainfall area, storm motion, slow moving boundaries, and training storms. We will look at each one to see how they can increase precipitation duration.

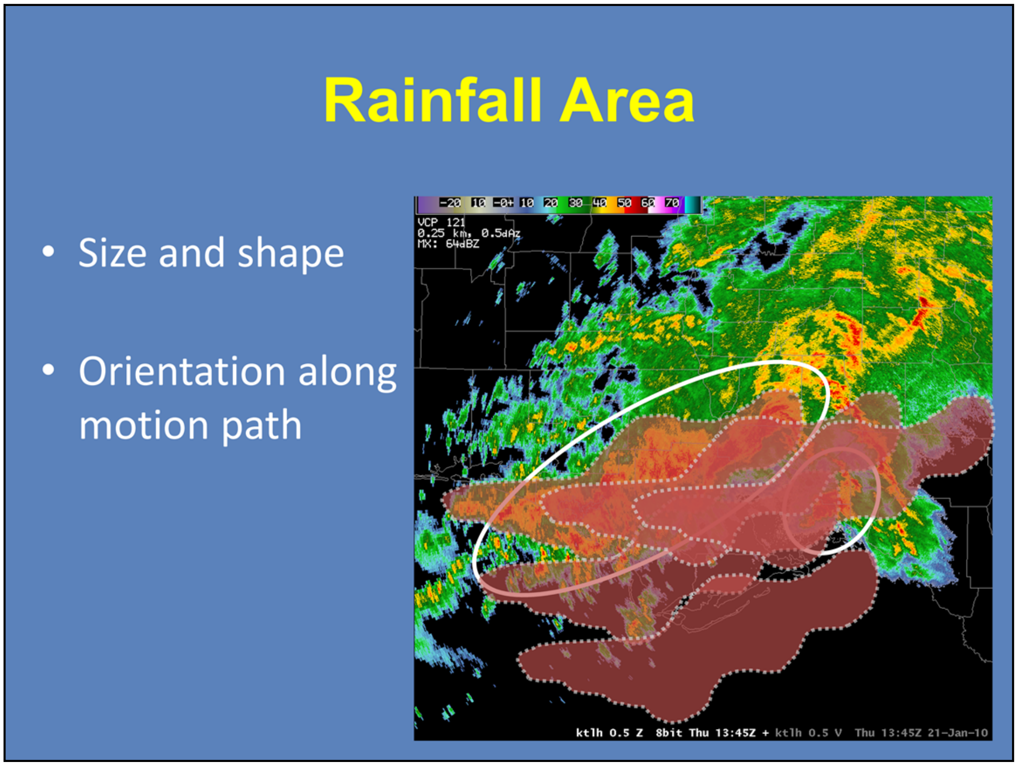

The first thing to look at is the size and shape of the precipitation area. Using the Tallahassee, FL radar, you can see a supercell southeast of the radar. Supercells and pulse storms are small in size, and depending on movement, will generally have a small residence time over one location. Linear convection, like the complex to the west, cover a much greater area. Therefore, the residence time of rainfall over a specific point is increased.

One thing to look at is the orientation of the precipitation area with respect to the motion path. Let’s assume that this convection highlighted here is moving towards the south at a constant speed of 40 mph. If we were to assume that the width of the area is approximately 20 miles, then the residence time of the moderate to heavy rain is about 30 minutes.

Now let’s assume that this linear complex is moving to the east at 40 mph. If we were to assume that the length of the area is about 120 miles, then the residence time over this area is closer to three hours. With this event back in January 2010, the complex was moving towards the east and produced 4-7 inches of rain around the Tallahassee area.

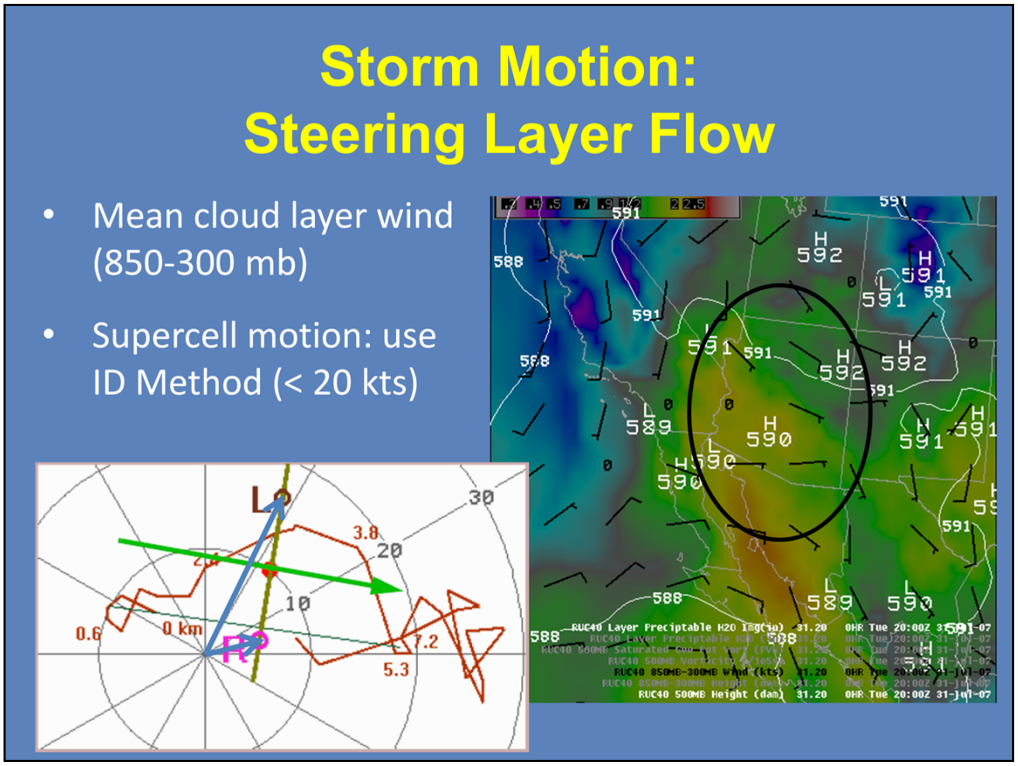

Storm motion is a significant factor when it comes to precipitation residence time over an area. Obviously, slower storm motions lead to longer durations. But what would you look for to determine storm motion?

One factor to look at is the steering layer flow. You can use your volume browser in AWIPS to view the mean wind between 850 and 300 mb. Using the example over Arizona, you can see that the 850-300 mb winds over the state is generally from east to west at 5 kts. Very slow moving storms in this area did produce fatal flash floods in the Tuscon CWA.

For supercells, you can use the Internal Dynamics (ID) Method to calculate the motion of right and left moving supercells. A storm motion of under 20 kts is preferred. In the example shown here, you can see that right moving supercells with this hodograph would be moving just north of due east at about 5 kts. Recall how to use the ID Method with hodograph in the lesson on Supercell Dynamics and Motion.

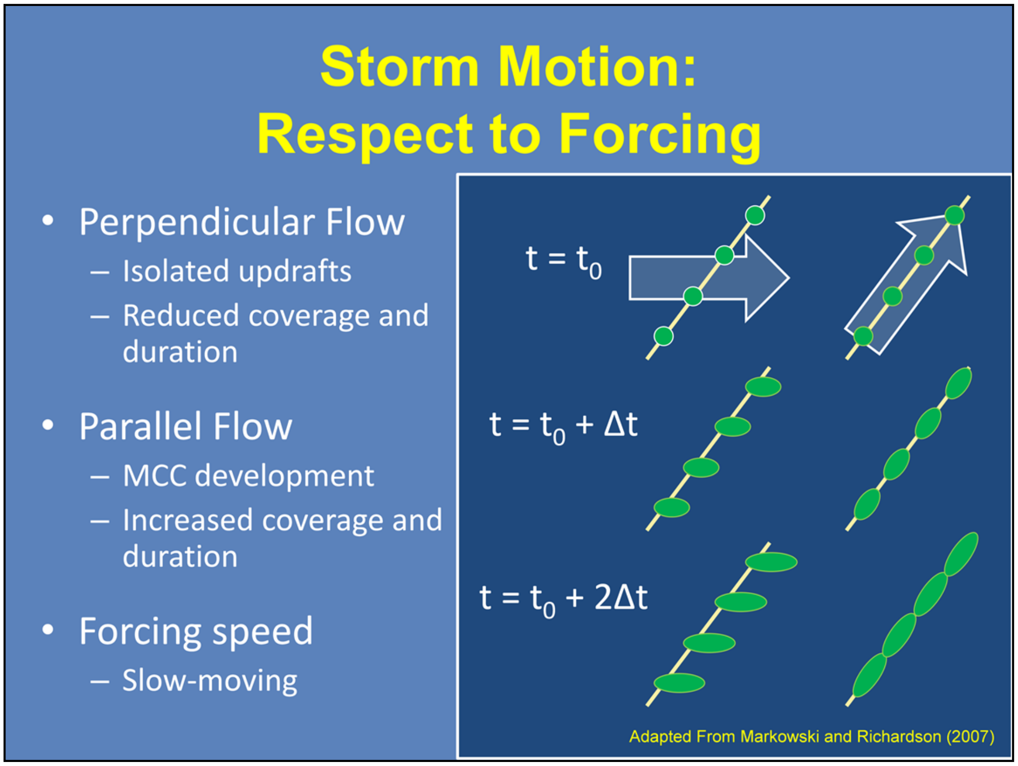

Forcing mechanisms play an integral role in the development and motion of convection. A forcing mechanism can range from fronts to outflow boundaries to topographic features. How storms form and move along a boundary can determine whether you have isolated updrafts or consolidated line segments and mesoscale convective complexes (MCCs).

Recall the work of Markowski and Richardson. Flow that is perpendicular to the forcing will lead to isolated updrafts, which in turn will have reduced areal precipitation coverage and smaller precipitation durations. Flow that is more parallel to the forcing will lead to linear convective formation. This will increase precipitation coverage and duration. Slow moving or quasi-stationary forcing mechanisms are best for increased precipitation residence time over an area.

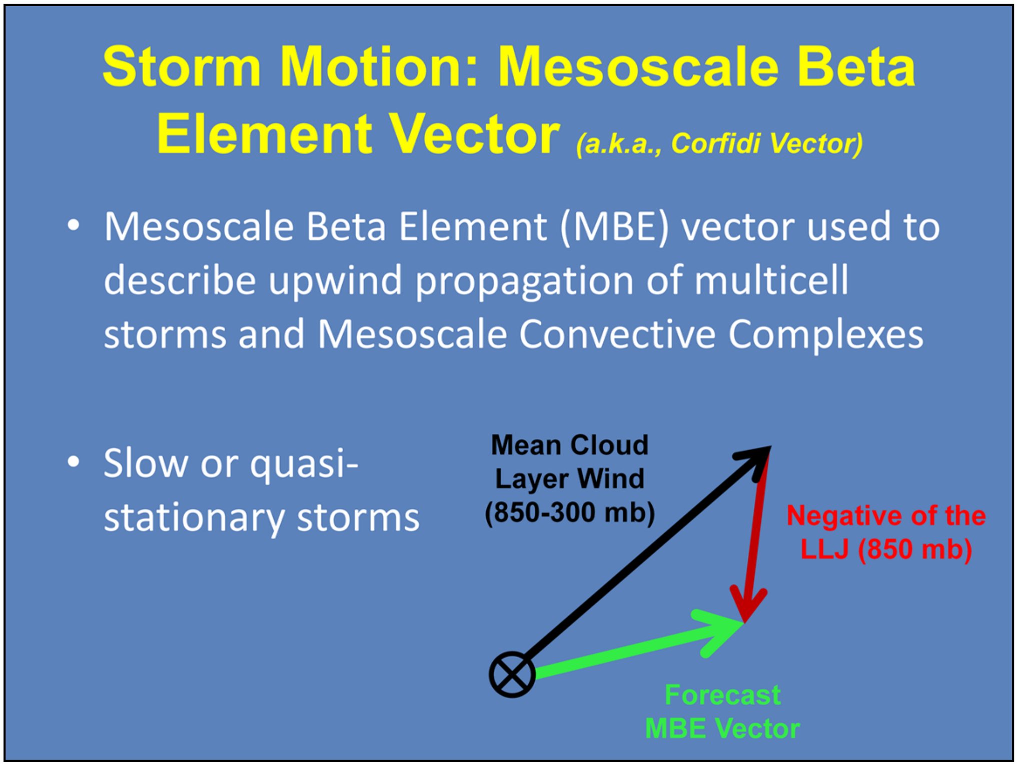

When dealing with multicell storms and mesoscale convective complexes (MCCs), the Mesoscale Beta Element (MBE) vector can help describe the upwind propagation of multicells and MCCs. Recall from the lesson on Multicell Motion that the MBE vector is calculated from taking the mean cloud layer wind and adding the negative of the low-level jet (850 mb flow depending on the depth of the inflow layer). Small MBE vectors means that if there is upwind propagation, then the complex will be slow moving or even quasi-stationary.

Back to Top

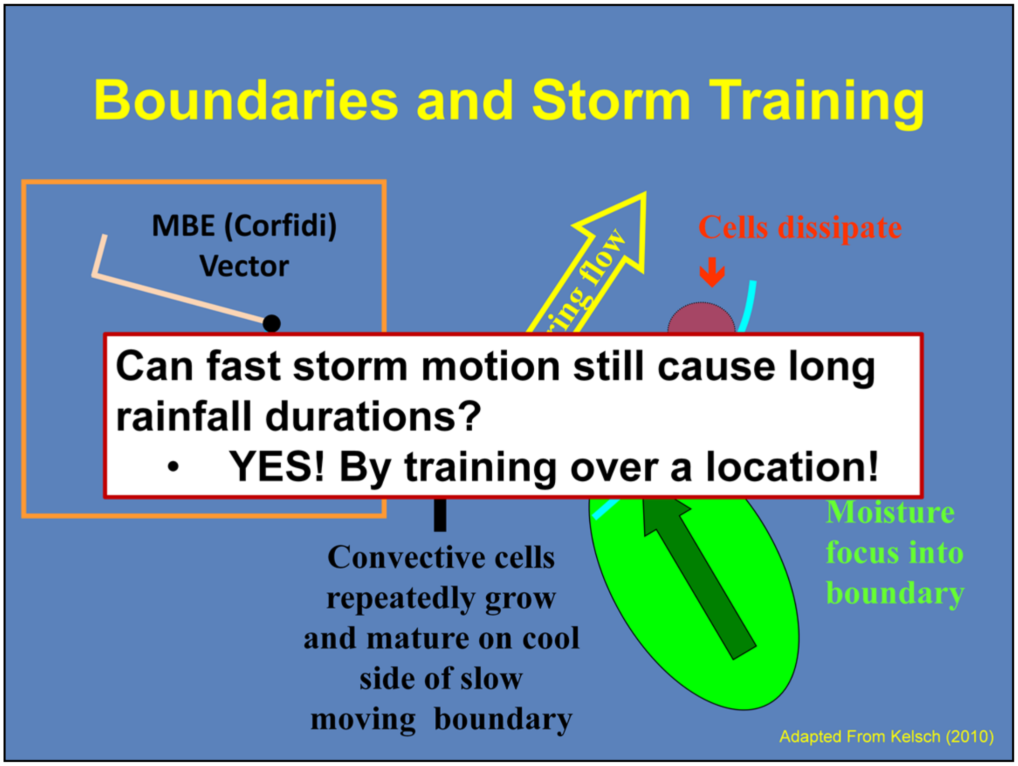

So far, we have talked about slow storm motions. What if storm motions are relatively fast? Can we still get large durations of rainfall? The answer is definitely yes.

If storms are training over the same location, it is easy to get the adequate duration for flash flooding to occur. One way is to have storms continuously propagate along a slow moving boundary. In this diagram, you have a SW-NE oriented boundary with an area of focused moisture transport. With enough lift and instability, convective cells will develop, move along the boundary, and dissipate. This cycle will continue so long as the boundary motion, moisture, instability, and trigger remain constant.

If you were to examine the vectors of this case, the mean flow parallels the boundary with expected storm motion of 25 kts. The MBE Vector shows that with backbuilding storms (upwind propagation), this system will move to the east at about 5 kts. This will allow for ample precipitation duration for flash flooding.

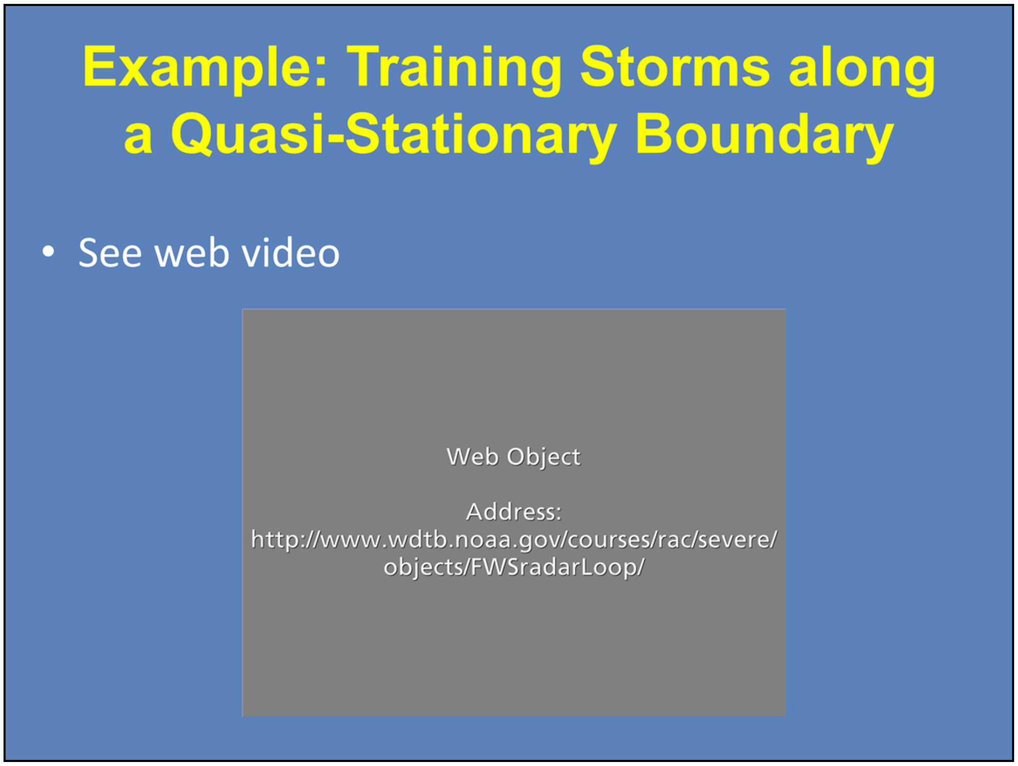

Here is an example of training storms that led to significant flash flooding. In this case from the Dallas/Fort Worth office, a series of storms train over the northern part of Texas near the Red River. This loop shows 5 ½ hours of radar data from KFWS. The star on the map shows the relative area of maximum focus and continuous development. Note how the storms train over the same area until a substantial cold pool develops for forward propagation.

Click next to advance to the analysis of this event when you are done viewing this loop.

Back to Top

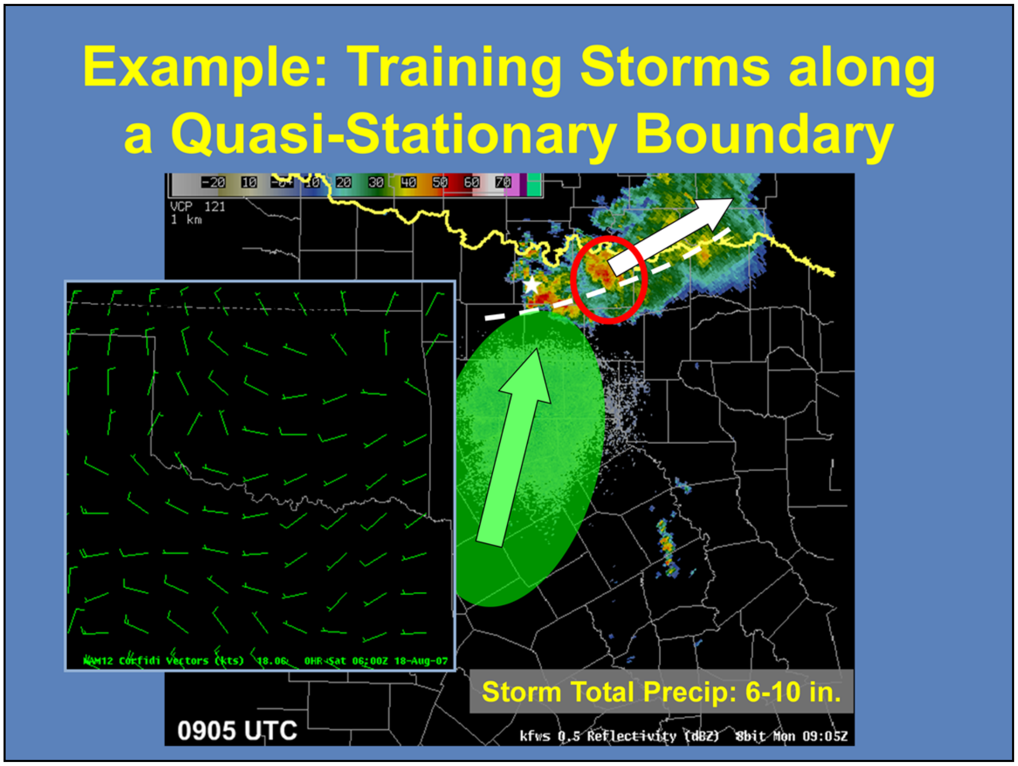

This event was created from a remnant mid-level circulation and boundary where a small vorticity maximum around the southern periphery is providing focus along the axis of forcing. The 1200 UTC sounding from Dallas/Fort Worth showed a very moist southerly 850 mb winds at 35 kts. You saw that storms initiated along the boundary where the forcing was maximized and then moved off to the ENE. However, the area of storms barely moved over a four hour time period. As you see here in the MBE, or Corfidi, Vectors, overall forecasted motion of the system is around 5 kts.

During this event, some areas received over four inches of rain in less than two hours, and storm total precipitation of 6-10 inches. There were six fatalities from these flash floods. Grayson County, which is circled in red here, had approximately 450 water rescues from vehicles and homes. There were hundreds of other water rescues in the surrounding counties.

Summary

Precipitation Rate

- Upward Moisture Flux
  - Updraft Strength (CAPE)
  - Vertical moisture profile
- Precipitation Efficiency
  - Warm vs. Cold Rain Processes
  - Warm Cloud Layer
  - Cloud Seeding

Precipitation Duration

- Precipitation area
- Storm motion and forcing mechanisms
  - Steering Layer Flow
  - MBE (Corfidi) Vectors
  - Quasi-stationary boundary
- Training Storms

In summary, you saw that there were two primary meteorological factors regarding rainfall and flash flooding. With precipitation rate, you saw how the strength of the updraft and the overall vertical moisture profile played a role in the upward moisture flux into a storm. The fraction of that that is returned as precipitation is defined as the precipitation efficiency of a storm. This is dependent upon the type of rain processes that are dominant, the depth of the warm cloud layer, and cloud seeding.

With precipitation duration, you have to consider the area and motion of the precipitation. Understanding storm motion and forcing characteristics, such as flow parallel to boundaries, weak steering layer flow, backbuilding complexes via slow MBE vectors, and slow or quasi-stationary boundaries, can help provide longer duration periods. Analyzing the mesoscale environment can help you determine the potential for training storms if storm motions are relatively fast.

Jill Hardy again and welcome to this lesson on flash flood hydrology. Steve Martinaitis of OU CIMMS at NSSL will again be narrating. But if you have any questions regarding the material, please feel free to contact me, or the RAC team. Our contact information will be on the next slide.

Back to Top

Lesson Objectives

- Identify the basic details of flash flood guidance
- Identify the hydrologic characteristics that impact the flash flood potential and flash flood guidance

There are two objectives to this lesson. At the end of this presentation, you should be able to identify the basic details regarding the creation of flash flood guidance and to identify the hydrologic characteristics that can impact the flash flood potential and the flash flood guidance product.

Back to Top

Why Hydrology is Important

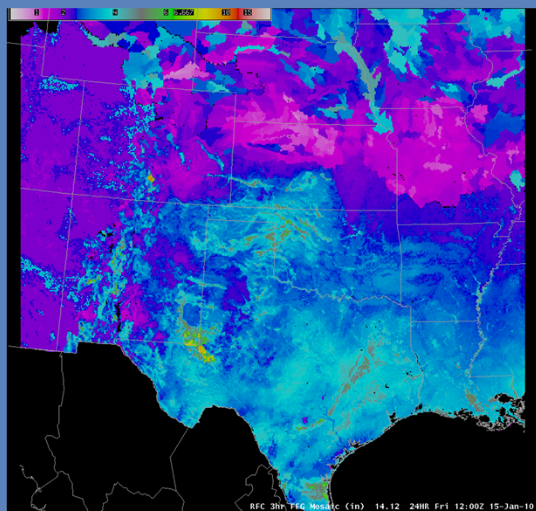

Flash Flood Guidance

Amount of rainfall needed within a certain period of time for small streams to overflow their banks

Understanding hydrologic properties, such as basin geometry, land use, and soil moisture, can help in determining what areas are prone to flash flooding and why. The flash flood guidance (FFG) product is created using these hydrologic properties to provide NWS forecasters a rainfall value needed within a certain temporal period for small streams and creeks to overflow their banks. And as you can see, these values can vary quite drastically across the country.

Back to Top

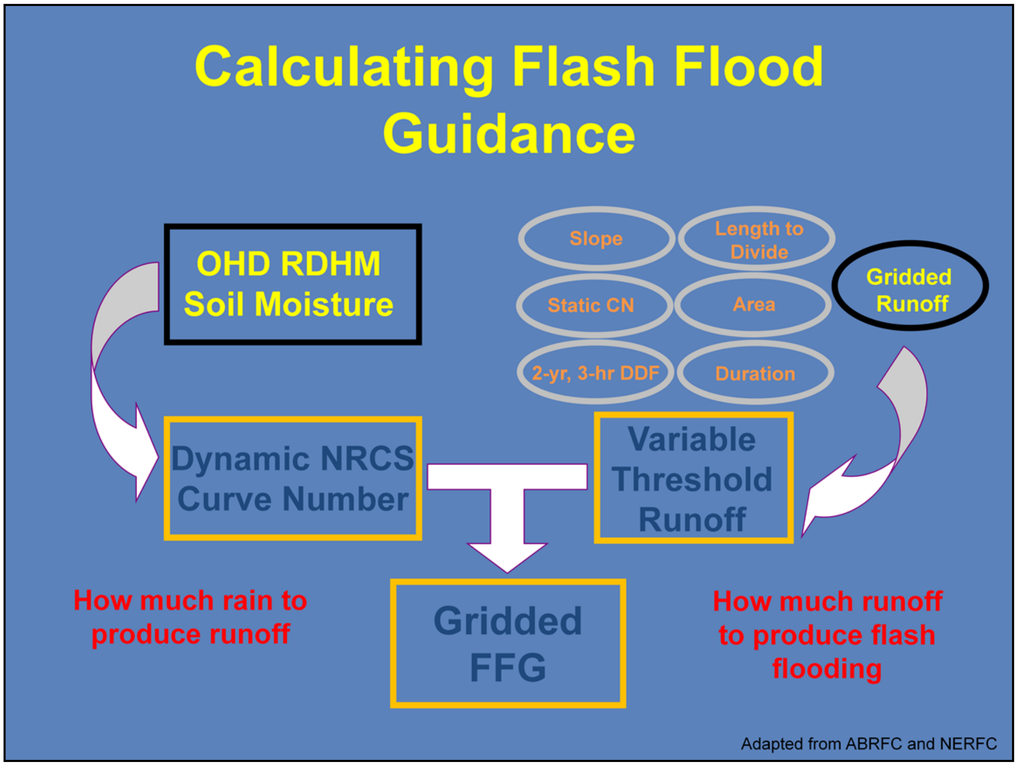

Most river forecast centers (RFCs) create a gridded flash flood guidance (FFG) product generally four or more times a day. FFG is derived from how much rainfall is needed to produce runoff via a dynamic National Resources Conservation Service, or NRCS, curve number (CN) and how much runoff is needed to produce flash flooding via a threshold runoff, or Thresh-R, value.

Both the NRCS curve numbers and Thresh-R values are calculated using different hydrologic properties. We will first focus on what hydrologic conditions influence the NRCS curve number.

Back to Top

Defining the Curve Number

- National Resources Conservation Service (NRCS) Curve Number (CN) – Empirical parameter for predicting direct runoff

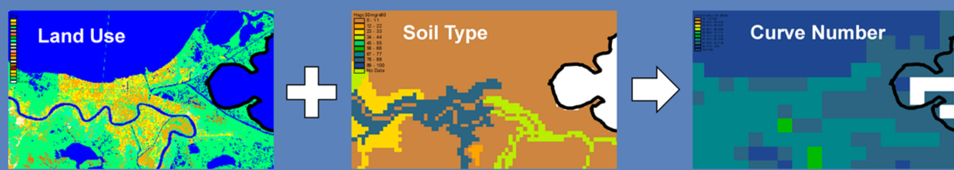

- Let's first talk about some of the land/soil properties…

The National Resources Conservation Service, or NRCS, Curve Number (CN) is an empirical parameter for predicting direct runoff. The curve number is generally based on soil type and land use, which both impact the amount of rainfall that is intercepted and infiltrated. The higher the curve number, the less rain is needed to create runoff. So, before we go into the operational details of the curve number, let's focus on the characteristics of land usage and soil types.

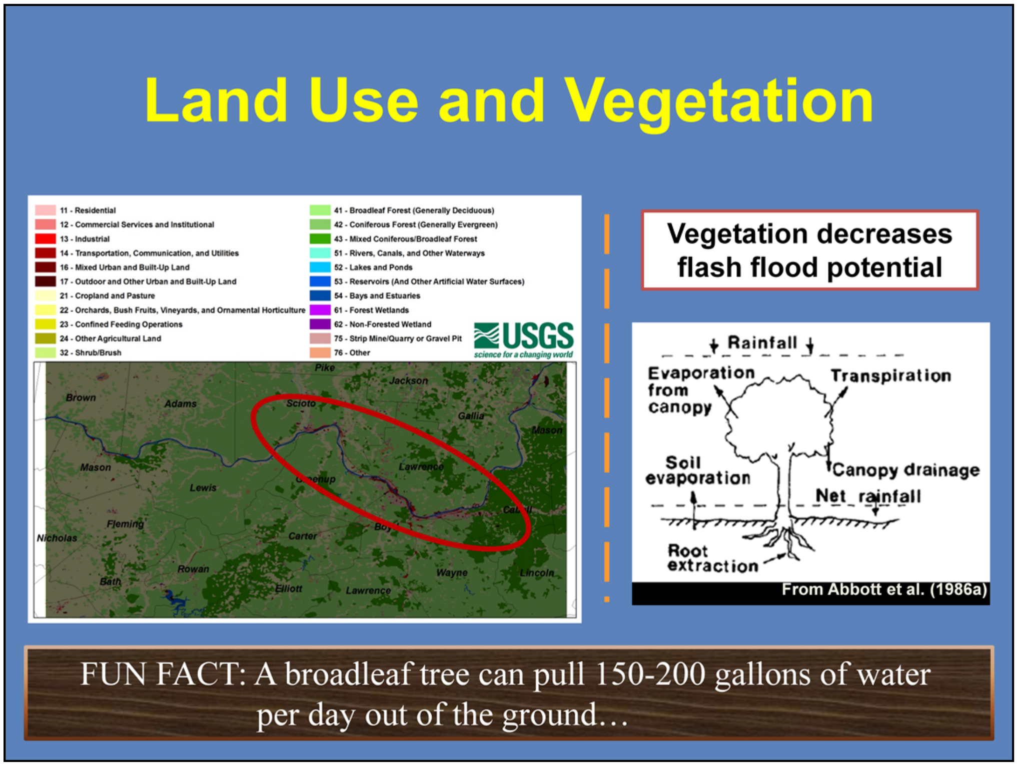

Land use can help determine how much water can be intercepted and how much water can be translated into runoff. Shown here is an example of a land use map for southern Ohio and northern Kentucky. The western part of this domain is dominated by croplands and pastures. The eastern two-thirds is a combination of broadleaf and coniferous forests. Urbanized areas are located along the Ohio River, and we will talk about the impacts of urbanization later in the module.

One the bigger influences of land use on the curve number is vegetation. Vegetation of all types help decrease the flash flood potential. Leaves can intercept rainfall before it reaches the surface. Water on and within the foliage undergo evapotranspiration. Roots help increase infiltration and can extract water from the surface and top layers of the soil. Did you know that one broadleaf tree can extract 150-200 gallons of water from the ground in one day? That’s a lot of water. This is why areas devoid of vegetation or have been deforested have higher curve numbers and a higher potential to flash flood.

Back to Top

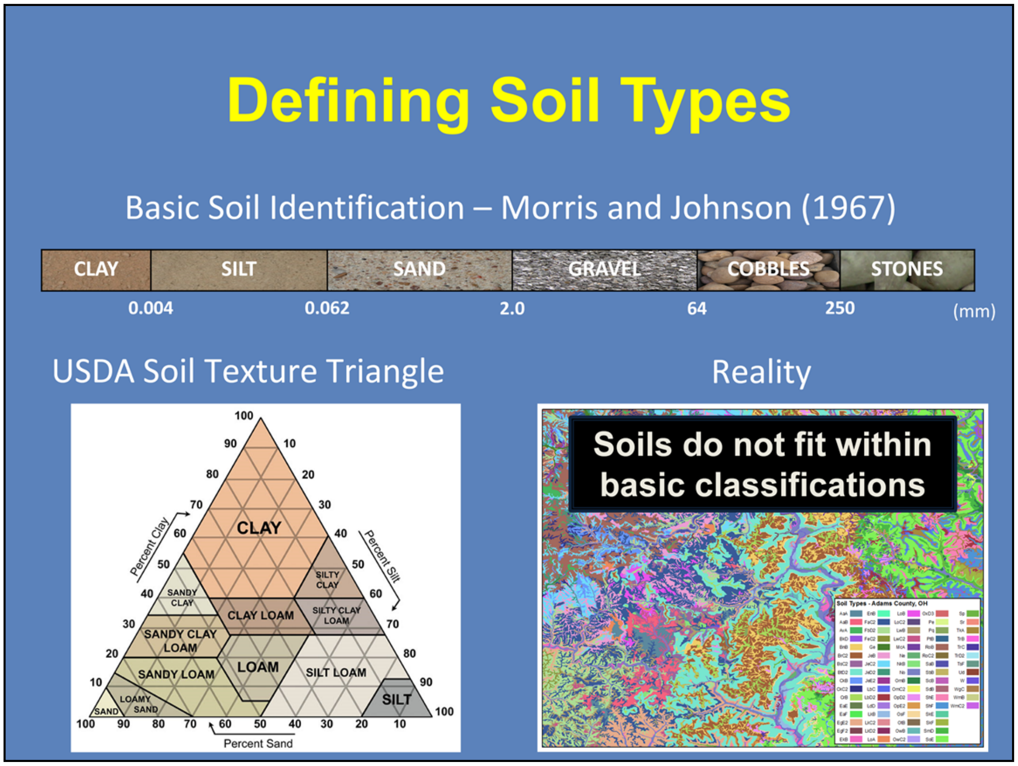

How soils are defined has evolved over the decades. The work by Morris and Johnson in 1967 defined the three basic soil types by the average size of their particles. Clay was defined by particles less than 0.004 mm while sand particles can approach 2 mm in size. The USDA then developed a soil texture triangle to better determine soil types based on the proportion of sand, silt, and clay after particles larger than sand have been removed. The USDA soil texture triangle defines 12 different soil types, including loam, which is a soil composed of sand, silt, and clay in relatively even proportion, and these types are commonly used in curve number and modeling calculations.

In reality, there are dozens upon dozens of soil types in the country. This image here shows 76 different soil types just within Adams County, OH. And in reality, soils do not fit within their basic classifications, such as those described by Morris and Johnson (1967).

Back to Top

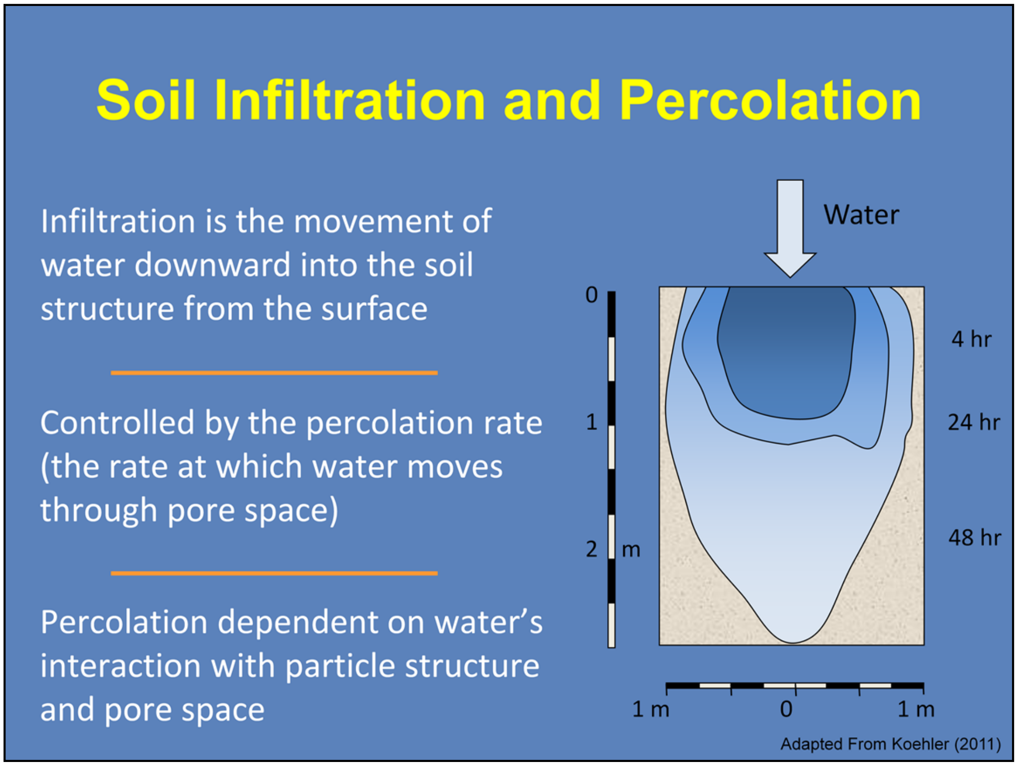

One of the most important characteristics of soil that impact runoff potential is the ability of water to infiltrate the soil. Infiltration is the downward movement of water from the soil surface into the soil structure. The ability of water to infiltrate the soil structure is controlled by the percolation rate, which is defined as the rate at which water moves through pore space. The percolation rate is generally dependent upon how water interacts with the particle structure and volume of porous space. For example, the pore space of clay-type soils can vary, but how water interacts with the clay particle reduce the infiltration and percolation rates. Thus, it takes less water to generate runoff. There are a number of different equations that can be used to calculate soil infiltration, which will not be discussed here.

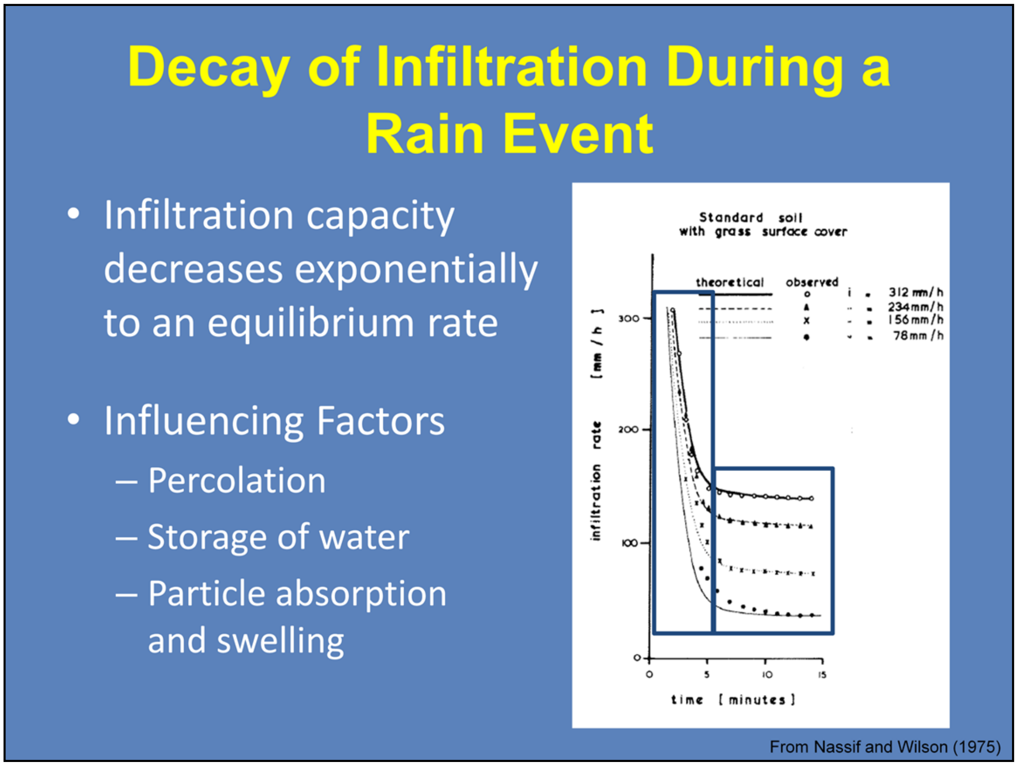

One important thing to know is that the maximum infiltration rate of the soil, more commonly referred to as the infiltration capacity, is *not constant*. A study by Nassif and Wilson (1975) compared rainfall rates with the infiltration properties of various soils. They tested different soil types, including some with grass surfaces, at different slopes using large soil trays and a sprinkler-type system to simulate instantaneous rain rates from 3-12 in./hr.

During their experiments, they found that within the first 5-10 minutes of applying the simulated rainfall, the infiltration rates of the soils decreased exponentially and began reaching an equilibrium rate. The experiment also demonstrated why a saturated ground cannot take in as much water as an unsaturated ground. The biggest factor that influences the ability of water to infiltrate the soil is percolation, while storage of water in the soil and particle absorption and swelling also influence this process.

Soil Infiltration and Curve Number

HSG	Soil Type	Characteristic
A	Sand, Loamy Sand, Sandy Loam	Low runoff potential, high infiltration rates
B	Silt Loam, Loam	Moderate infiltration rates
C	Sandy Clay Loam	Low infiltration rates
D	Clay Loam, Silty Clay Loam, Silty Clay, Clay	High runoff potential, low infiltration rates

- Hydrologic Soil Group (HSG) classification used to determine NRCS Curve Number

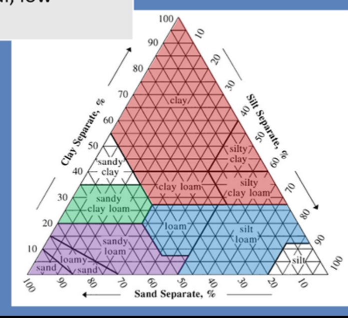

Knowing all of this, we look back at how soil types and land use are operationally applied. NRCS soil scientists categorized the soil types defined by the USDA soil texture triangle into four Hydrologic Soil Groups (HSGs) based on infiltration, runoff characteristics, and texture. Group A consists of mostly sandy-based soils that have low runoff potential and high infiltration rates. Group B consists of silt loam and loam that have moderate infiltration rates and moderately coarse to fine particle textures. Group C consists of sandy clay loam, which is has low infiltration rates with moderately fine textures. And finally, Group D consists of mostly clay-based soils that have high runoff potential and low infiltration rates.

These four HSGs are then used to determine the NRCS curve number.

Back to Top

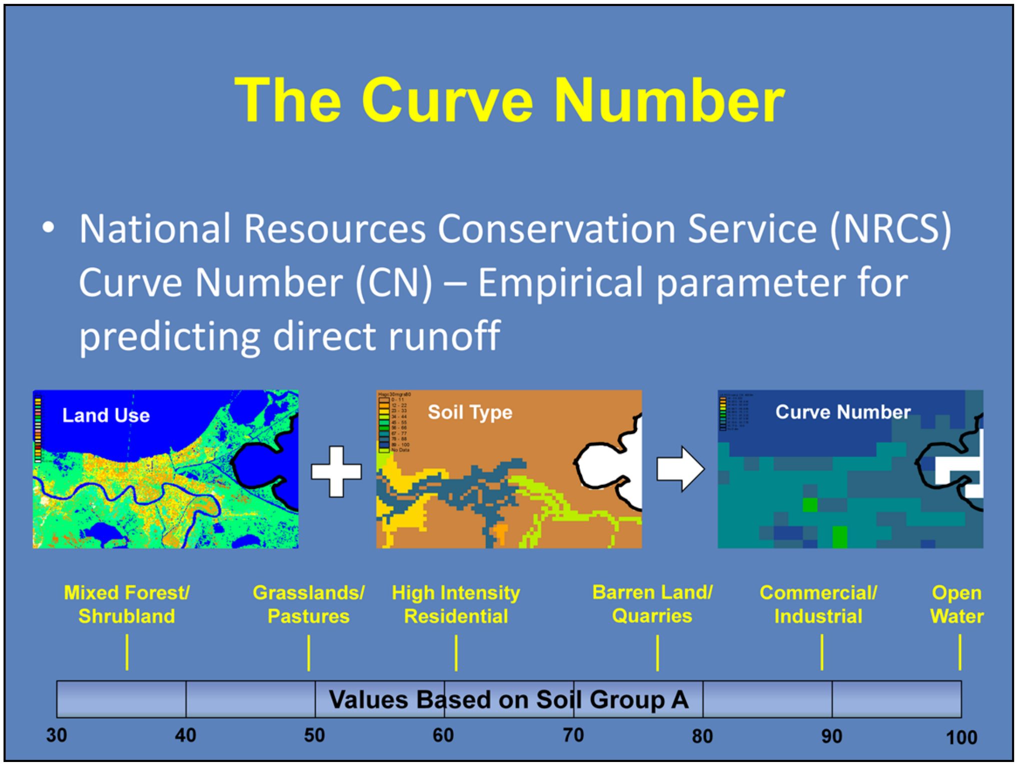

So let's see how the land use and soil types impact the NRCS curve number (CN). The CN is a unit-less number that ranges from 30 to 100. The higher the number, the less rain is needed to create runoff. Each type of land use is given four CNs, one for each Hydrologic Soil Group (HSG). In this example here, we will define the curve number for a variety of land uses for soil Group A, the sandy-based soils with low runoff potential and high infiltration rates.

Mixed forests, shrublands, grasslands, and pastures all have low curve numbers, which means these land uses for this soil group have low runoff potential and would need a lot of rain to generate any runoff. Here are where high intensity residential areas, barren lands, quarries, and commercial/industrial areas fall on the curve number spectrum. Notice how more urbanized land and land devoid of vegetation have much higher curve numbers. Open water has a curve number of 100, meaning all rainfall becomes runoff. Remember that these values are for Group A, the sandy-based soils. The curve number for each land use type is much higher with Groups B, C, and especially D.

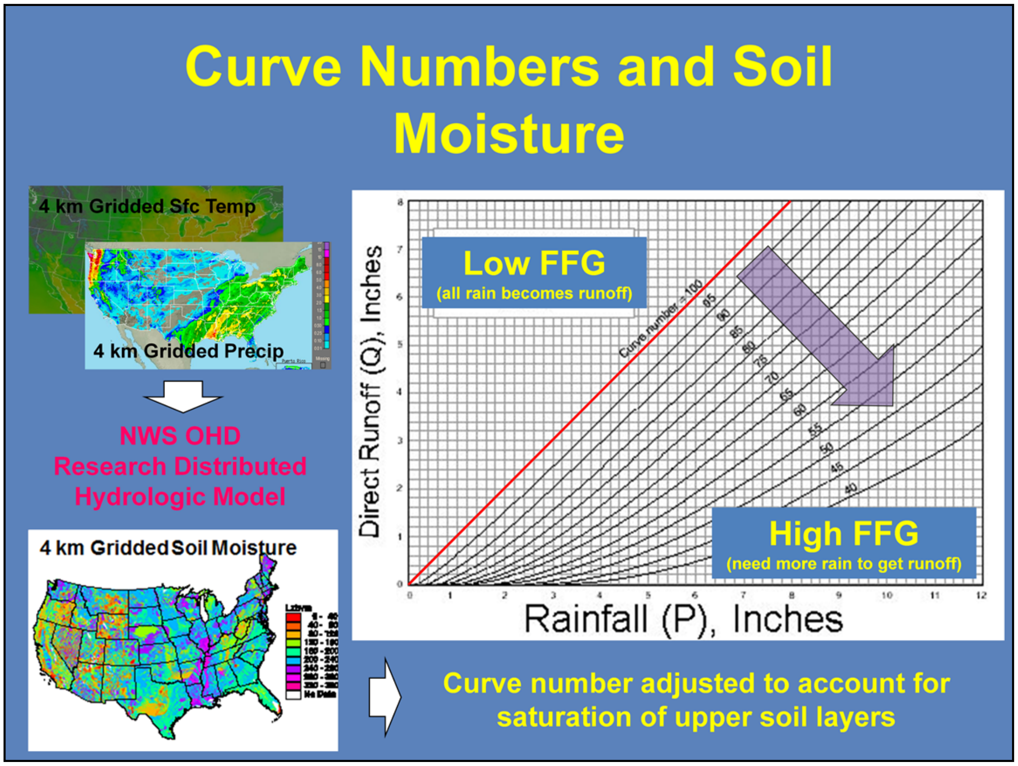

For gridded FFG, the National Weather Service (NWS) Office of Hydrologic Development (OHD) uses a research distributed hydrologic model to create a dynamic curve number. The model uses CONUS-scale 4 km gridded surface temperatures and 4 km gridded precipitation to create a gridded soil moisture product. The CN is then adjusted based on the soil moisture calculated to account for saturation of the upper soil layers. In areas that have seen recent, heavy rainfall, the curve number can be increased to near 100, which means that the FFG is reduced and nearly all rainfall can potentially be translated into runoff. The lower the curve number, the greater the FFG and the more rain that is needed to start generating runoff.

Back to Top

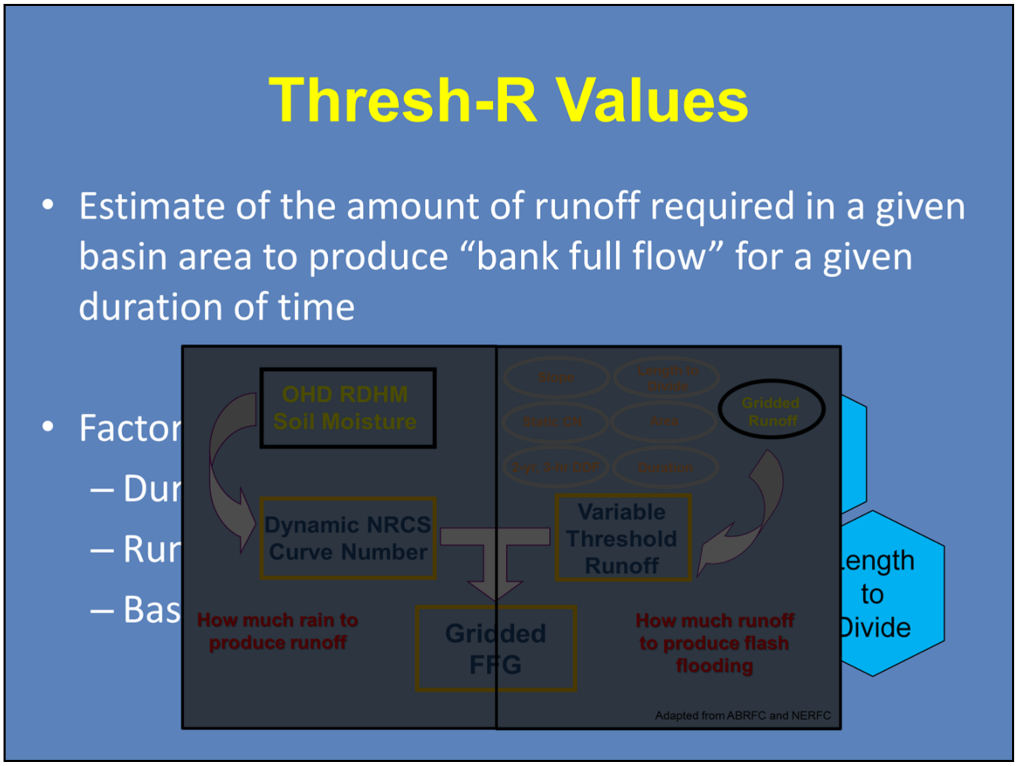

So we looked at the first half of the FFG calculation that asks how much rain does it take to generate runoff. Now we will focus on how much runoff will it take for small creeks and streams to reach "bank full flow" conditions, which is defined by the Thresh-R value.

The Thresh-R value considers a number of different factors, including duration of rainfall, runoff estimation, and a variety of basin characteristics. In this module, we will focus on what impacts basins have on the FFG.

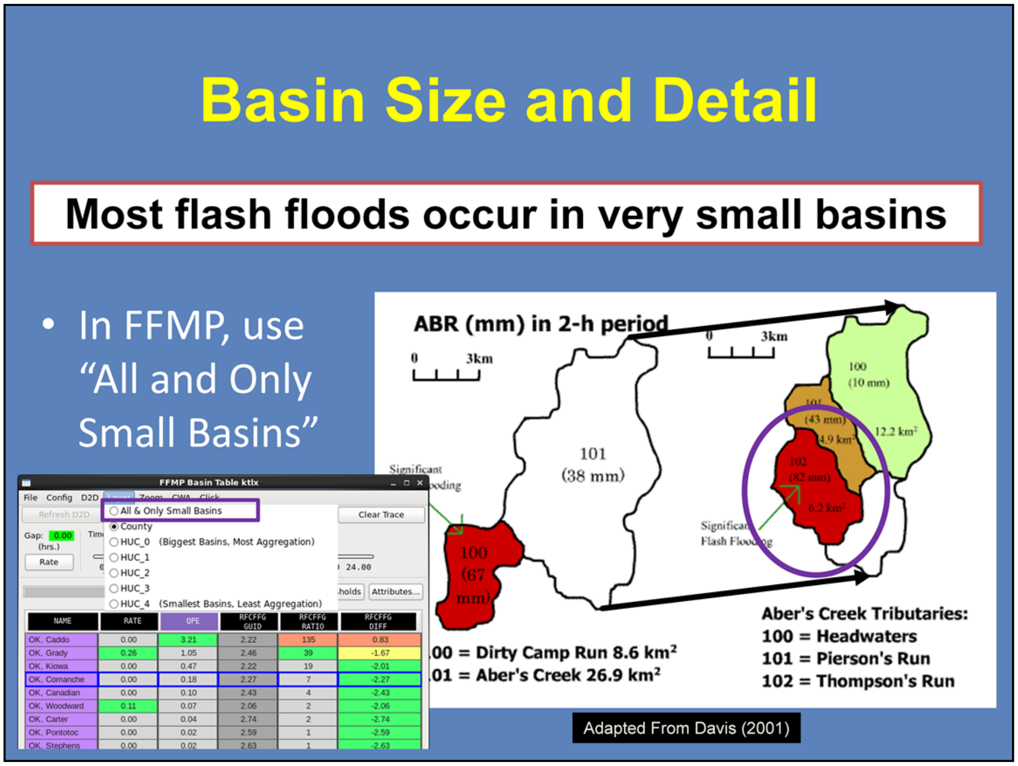

There are many levels of basin detail, ranging from the parent basin, such as the Mississippi River Basin or the Florida Watershed, to the smaller basins and sub-basins that compose them (a couple of square kilometers). The majority of flash floods occur in very small basins, mainly because the scale of the heaviest rainfall is also quite small.

Using this example from the Aberdeen, SD office, you can see two adjoining basins, Dirty Camp Run (#100) and Aber's Creek (#101). Significant flash flooding occurred in the Dirty Camp Run basin as a result of an average of over 2.5 inches of rain in two hours across the basin. In Aber's Creek, there was an average of 1.5 inches of rain across the basin. However, if you were to split the Aber's Creek basin into its sub-basins, you can see that the southwest part of the basin had almost 3.25 inches of rain in this two-hour period. Without looking at the smallest basins, it is possible that some areas would have went unwarned due to large scale basin averaging.

When using FFMP, you will get the greatest detail by going to the Layer menu in the FFMP Basin Table and selecting "All & Only Small Basins."

Back to Top

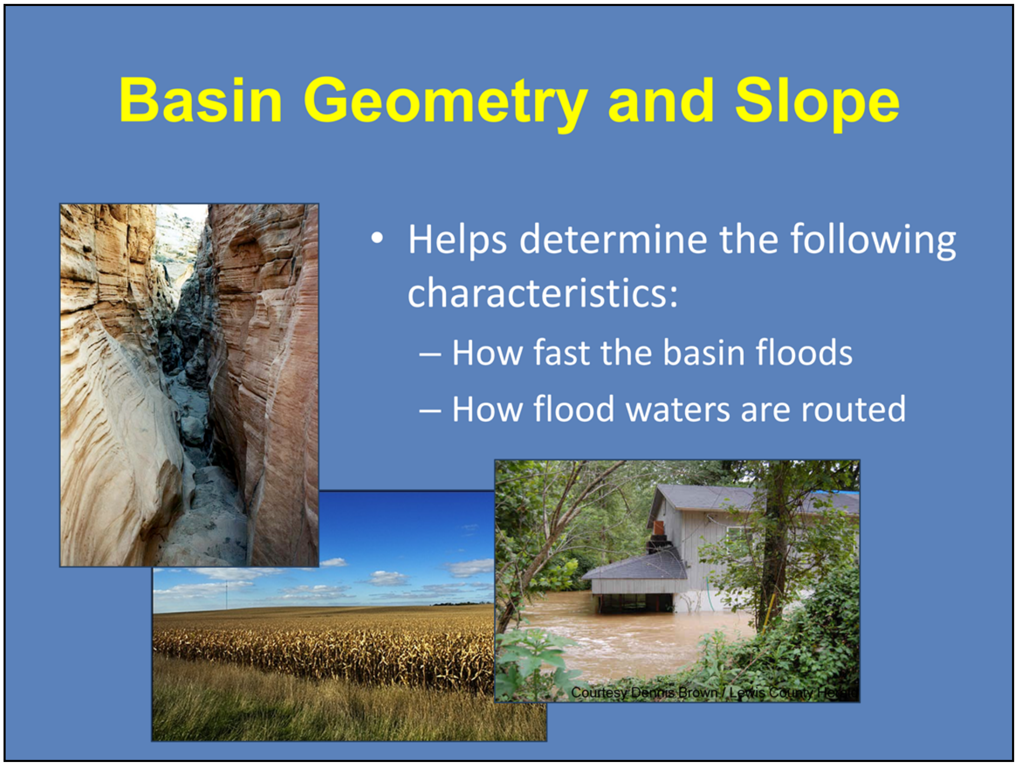

The three-dimensional geometric characteristics of a basin is important when determining how fast a basin can flood and how the flood waters are routed within the basin. A lot of this is dependent upon the slope of the terrain. If we are talking about the flat fields of the central Plains, then it could take awhile for runoff to move out of the area. In contrast, if you look at the mountainous terrain and slot canyons of the western U.S., then water can be quickly routed downstream. Think of a marble on a shelf. If the shelf is level, the marble is stationary. If the shelf sits at an angle, then marble rolls off. The same theory applies here.

Lag Time and Travel Time

- **Travel Time** – Time it takes a raindrop to travel between any two locations
- **Lag Time** – Time from the center of mass of excess rainfall to the hydrograph peak

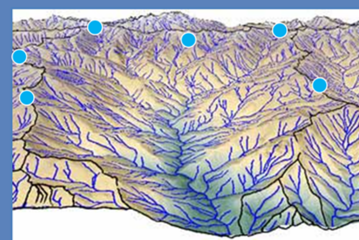

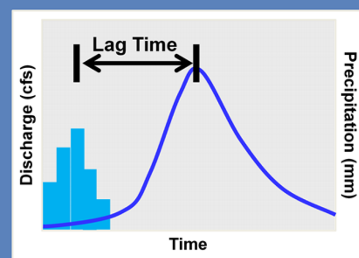

Two terms that are used when talking about runoff in a basin is lag time and travel time. Travel time is the time that it takes a raindrop of water to travel between any two locations.

Lag time is the time difference between the center of mass of excess rainfall and the peak in a river hydrograph.

Both of these are dependent upon the size and shape of the basin, as well as the slope of the ground within the basin. Basins that are prone to flash flooding can have a lag time of just 15 minutes and the water can sometimes travel a considerable distance in a short amount of time. Basins not prone to flash flooding can have lag times that can exceed one hour.

Back to Top

End Result = FFG

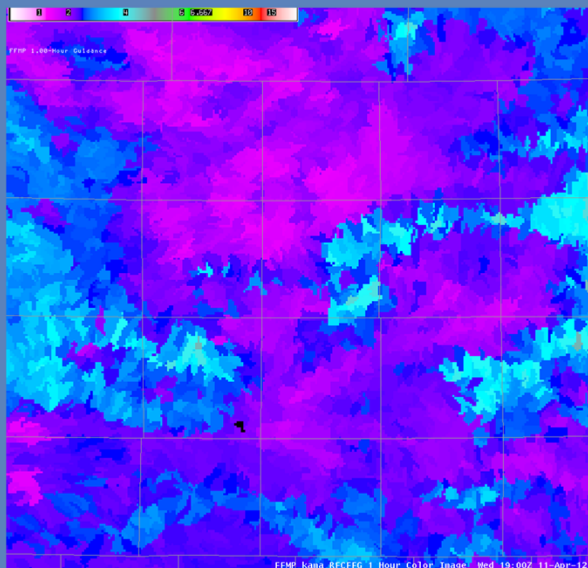

- Issued by RFCs generally four times a day
  - Updated more frequently as needed
- Adjusted to account for saturated soils during events
- Local forcing of FFG

Through hydrologic modeling, the combination of dynamic NRCS curve numbers and Thresh-R values creates the operational FFG that we use today. FFG is issued by the RFCs generally four times a day, and they can be updated more frequently as needed. Since gridded FFG accounts for recent changes in soil moisture, they can be adjusted for saturated conditions during events. FFG can also be forced locally using a GUI in the AWIPS workstation.

Areas that have compromised FFG, and where FFG is generally forced, are urban areas and wildfire burn scars. Let's start with how urbanization impacts FFG.

Urbanization can create as much as five times more runoff than that of a completely forested area. This primarily has to do with areas being covered by concrete instead of soils and vegetation. The infiltration rate of concrete is near zero, so almost all rain that falls on to concrete surfaces is translated into runoff. In addition to the “concrete jungle,” drainage systems within urban areas may not be adequate enough to handle very high quantities of water. But there is more to urbanization than just creating more precipitation runoff.

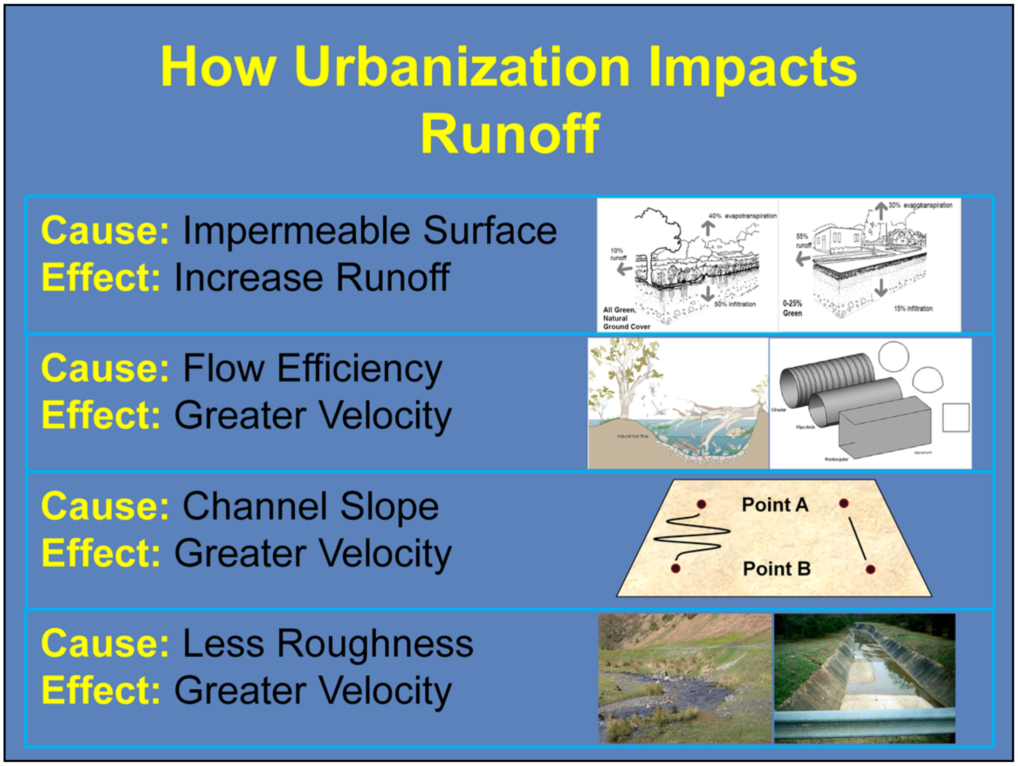

We know that when you have more impermeable surfaces (i.e., areas covered by concrete and asphalt), you are able to generate more runoff. But we must also look at what other factors urbanization can have on runoff.

Natural rivers and streams have very complex shapes, while urban drainage systems, culverts, pipes, etc. have simple shapes. This allows for greater flow efficiency of the runoff, and thus, gives the runoff a greater velocity.

Streams and rivers have a tendency to meander and are never really straight. Urban areas tend to have very straight drainage paths. When placing a meandering path and straight path on a sloped surface, not only does the straight path take less time for runoff to go from Point A to Point B, the overall slope of the straight channel is much greater. Straighter channels in sloped urban areas will have a greater channel slope and a greater runoff velocity.

Also, natural streams and rivers have a lot of roughness that come from rocks, vegetation, and complex channel shapes. On the other hand, urban channels, especially those that are paved, tend to be quite smooth, and thus have less roughness. This again translates into greater runoff velocity.

So, not only does urbanization create more runoff, it can also provide a greater velocity to the runoff, which makes it that much more dangerous to life and property.

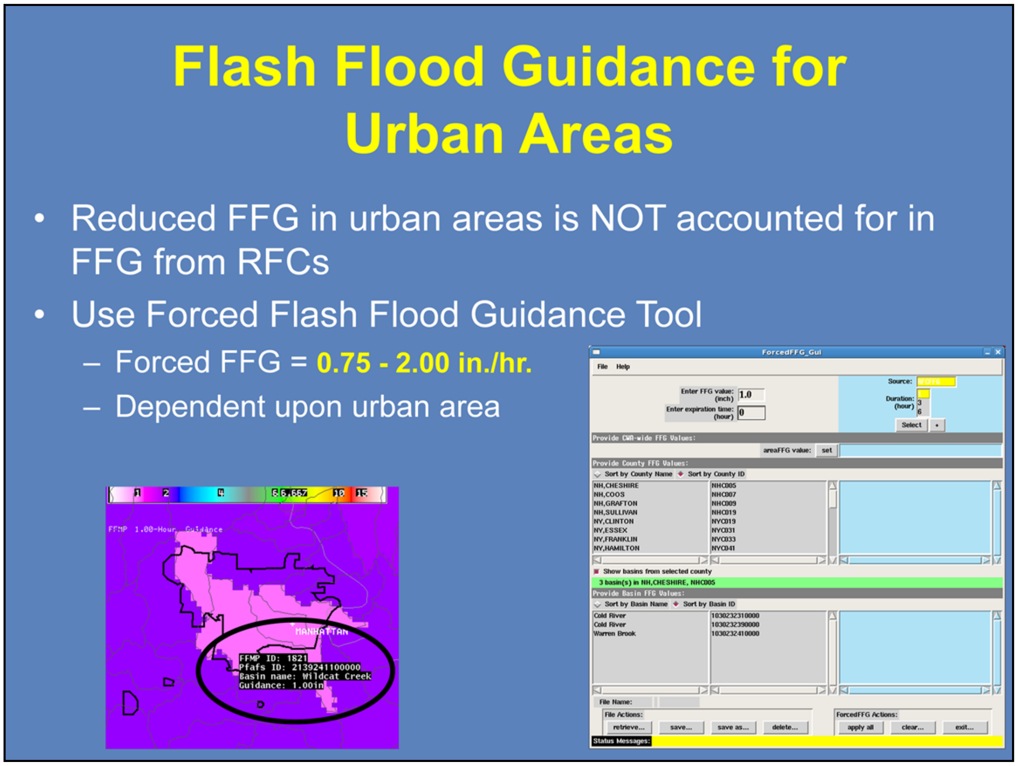

The reduced flash flood guidance (FFG) for urban areas are not accounted for in the FFG delivered by the RFCs. However, you can use AWIPS to force FFG for these areas. Depending upon the urban area, it can take anywhere from 0.75-2.00 in./hr. for flash flooding to occur. Ask your AWIPS focal point about using the Forced FFG interface.

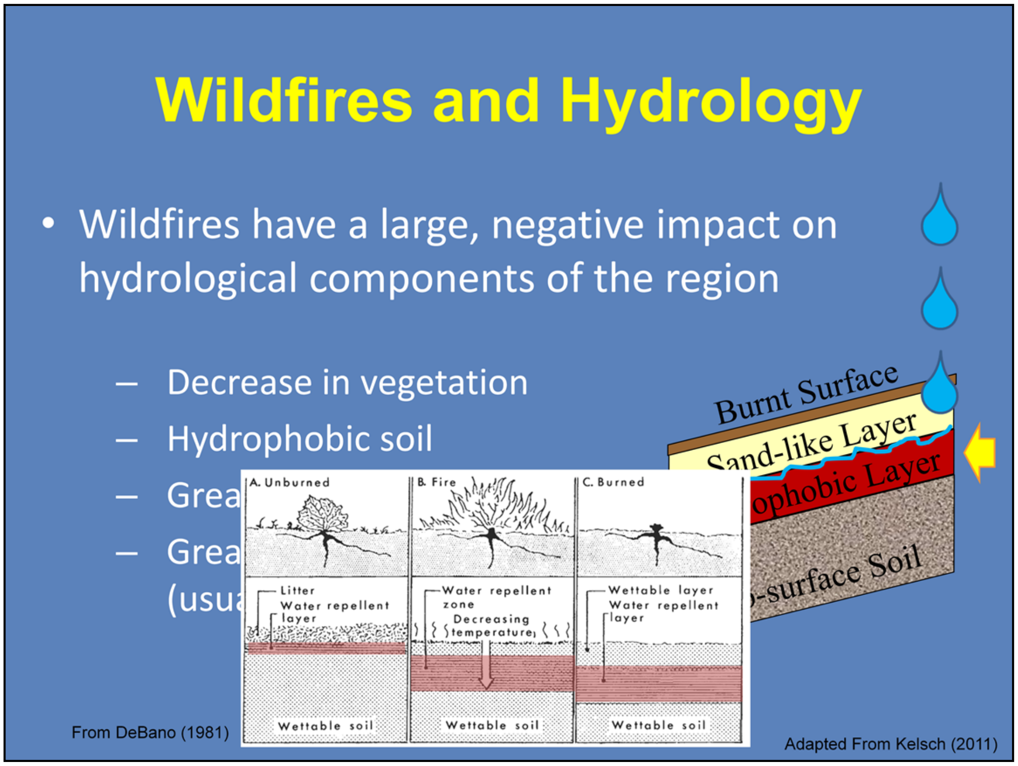

Now, I'm going to make a pretty obvious statement here and say that wildfires have a large and very negative impact on the hydrologic factors related to flash flooding. But, it is how the fires impact the area and for how long those impacts last that are critical to hydrology and flash flooding.

The part that we do see is the removal of vegetation. As we discussed earlier, vegetation helps capture water and increase soil infiltration, which reduces the potential for flash flooding. Without vegetation, there are no leaves to intercept rainfall or roots to extract water. One of the byproducts of combusting vegetation, especially in high-intensity forest fires, is the creation of a heavy gas that sinks and penetrates the soil profile. As this gas cools, it condenses and solidifies into a hydrophobic waxy coating around the soil particles. This is the part that we don't see and is very critical to flash flooding in burn scars. So, in the aftermath of high-intensity fires, a water-repellent sub-surface layer of soil is present. The greatest impacts from this occur within the first year after the fire, and these impacts remain for three to five years afterwards.

So, when it rains, the water will penetrate an initial layer of burnt soil and surface material and then a less-stable sand-like layer (all byproducts of the fire). When the water reaches the hydrophobic layer, it cannot penetrate any further. The water then becomes runoff with a high yield of sediment from the soils above this hydrophobic layer. Flash flood guidance is greatly compromised in burn scars, and it is usually around 0.50 in./hr.

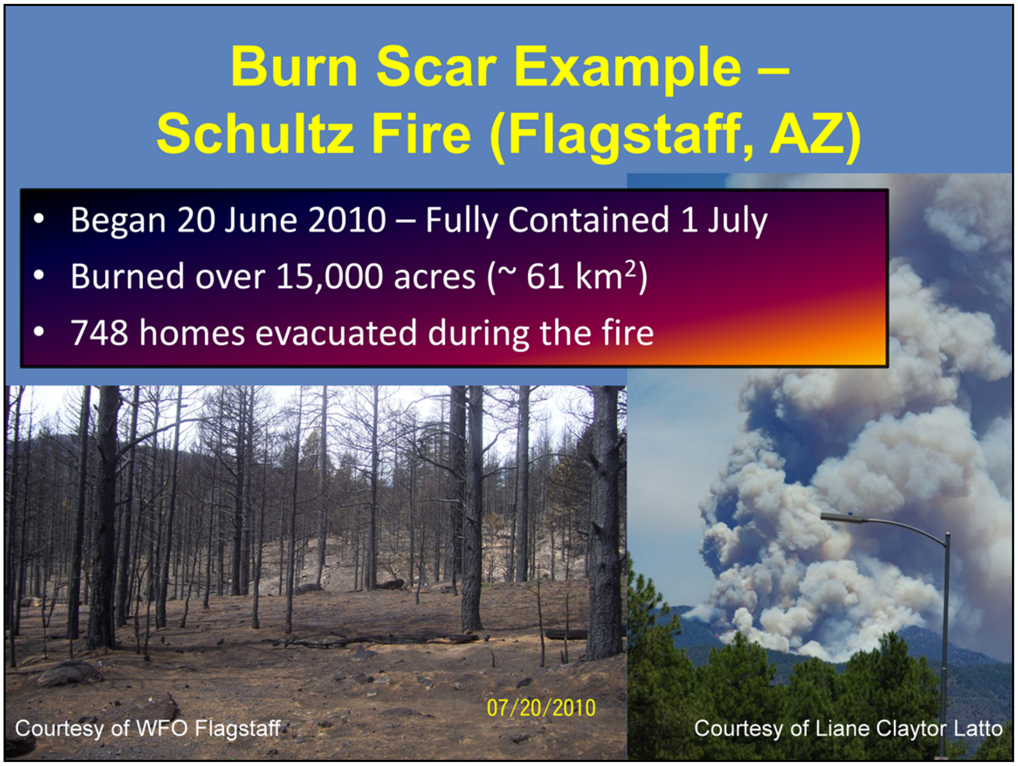

To show the impacts of a heavy rainfall event over a burn scar, we will look at the Schultz Fire of 2010. The fire began on the west side of the San Francisco Peaks, a volcanic mountain range north of Flagstaff, AZ, on June 20th, and the fire lasted about 12 days. The fire burned approximately 15,000 acres and resulted in the evacuations of nearly 750 homes.

The picture on the left shows what the terrain looks like one month after the fires. Note how the trees are devoid of any green vegetation and how the ground is burnt and lacks any underbrush or grass. Two hours after when this picture was taken, convection began to develop in the area.

Burn Scar Example – Schultz Fire (Flagstaff, AZ)

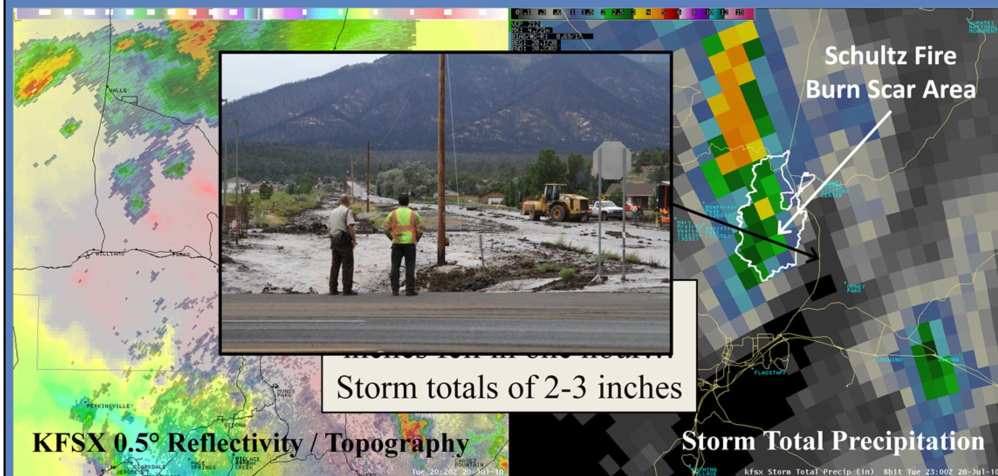

- Hundreds of homes with property damage, impassable roads, and one fatality

On the left is the KFSX 0.5° reflectivity and topography image combination. The blue contour north of Flagstaff represents the burn scar area from the Schultz Fire. On this day, there were slow moving storms in and around the mountain range. Once convection developed within the burn scar, the office issued a Flash Flood Warning in anticipation of flash flooding based on very slow storm motions and high precipitation rates.

The image on the right shows the storm total precipitation from the event. Two to three inches of rain fell within the burn scar, most within a one hour time period. The runoff and sediment quickly moved down the mountain, especially with the aid of paved roads, including Highway 89. Downstream suburban areas were heavily impacted with home and property damage and impassable roads. A 12 year old girl was killed during the event when she was swept away by the flood waters in her neighborhood. This was a case where you had runoff easily generated in a burn scar area, which quickly moved down a steep mountainside and traveled into an urbanized area where roads and culverts quickly channeled the water in this dangerous situation.

Back to Top

Summary

- Basics of Flash Flood Guidance (FFG)
- Hydrologic Factors that Influence FFG
  - Land Use and Vegetation
  - Soil Infiltration and Soil Moisture
  - 3-D Basin Characteristics
- Areas with Compromised FFG
  - Urbanization
  - Burn Scars

In summary, you should have a basic understanding what Flash Flood Guidance (FFG) is and how it is calculated at the RFCs. We looked at a number of hydrologic factors that influence the values of FFG. We saw how land use, vegetation, and soil properties influence the dynamic curve number, which describes how much rain is needed to create runoff. We also saw how different basin characteristics help influence the Thresh-R value used to determine how much runoff is needed to create bank full conditions.

Finally, we took a look at two areas with greatly compromised flash flood guidance: Urbanized areas and burn scars from wildfires. The impacts of urbanization, such as large areas of concrete and hydraulically efficient channels, yield greater and faster moving water runoff. Meanwhile, the impacts of wildfire burn scars can last for years due to the removal of vegetation and the creation of a hydrophobic layer in the soil. Overall, both lead to situations where only a moderate amount of rain can produce dangerous flash flooding situations.

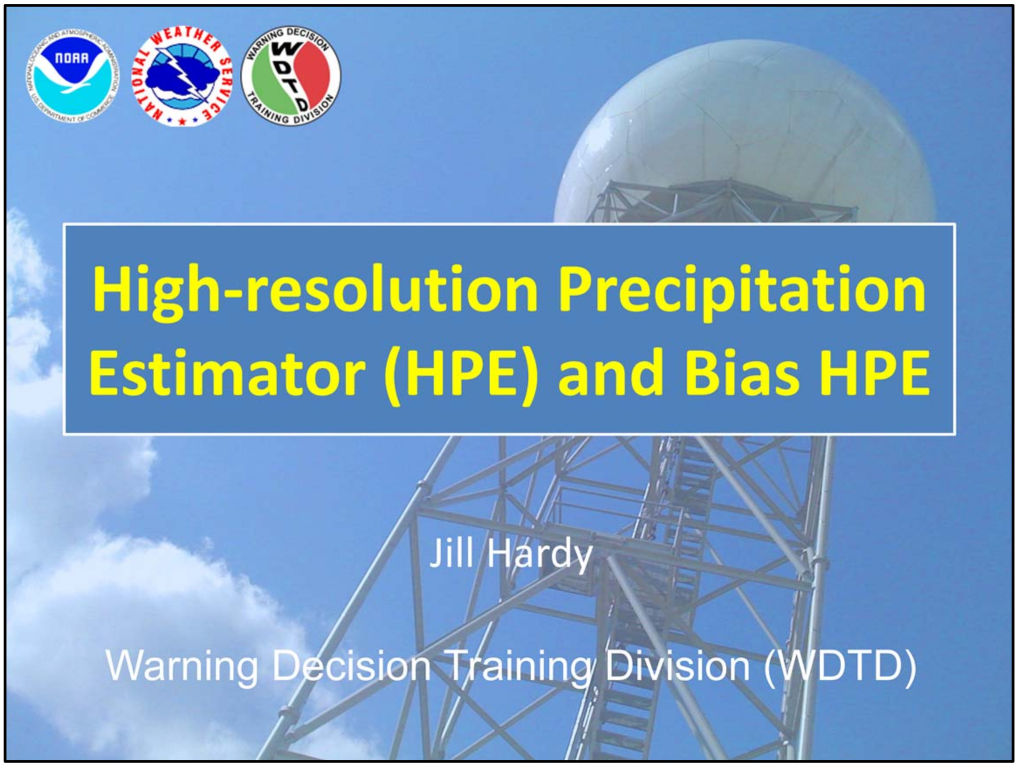

Hi, my name is Jill Hardy and welcome to this lesson on the High-Resolution Precipitation Estimator (or HPE) and Bias HPE products. Let’s jump right in!

Back to Top

Learning Objectives

- By the end of this lesson, you will be able to:
  - Define what are the HPE, Bias HPE, and HPN products
  - Identify the default precipitation source for HPE and Bias HPE
  - Identify the default bias source for Bias HPE
  - Interpret HPE and Bias HPE text overlays
  - Identify how to determine areas of higher confidence in HPE/Bias HPE precipitation estimates

Here are the learning objectives for this lesson. When you have finished reading them, please move onto the next slide.

What is the High-resolution Precipitation Estimator (HPE)?

Mosaic of rainfall data from all radars within your coverage area

Uses 0.5° scan
- Usually nearest radar
- Radar outage: higher altitude 0.5° estimates

HPE Products
- Instantaneous rain rate
- Accumulated precip (Default = 60 min)

Resolution
- 1-km by 1-km
- 5 min

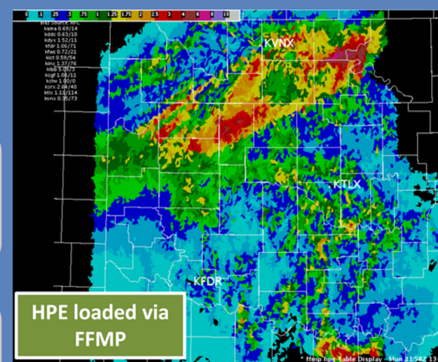

--The High-Resolution Precipitation Estimator (or HPE) is a mosaic of rainfall data from all radars within your coverage area.
--The mosaic is created using the lowest altitude scan, which is usually the nearest radar. However, when a radar is down, the mosaic will employ neighboring radars which will result in higher altitude estimates.
--The two available HPE products are the instantaneous rain rate field and the accumulated precipitation field.
--A one-hour accumulation is the default in HPE, but this value can be altered with help from your hydro focal point.
--These products update every 5 minutes. However, the one-hour product needs at least an hour's worth of data before it begins updating every 5 min.
--HPE will stop collecting data once every contributing radar has not received rainfall within the last 20 min.

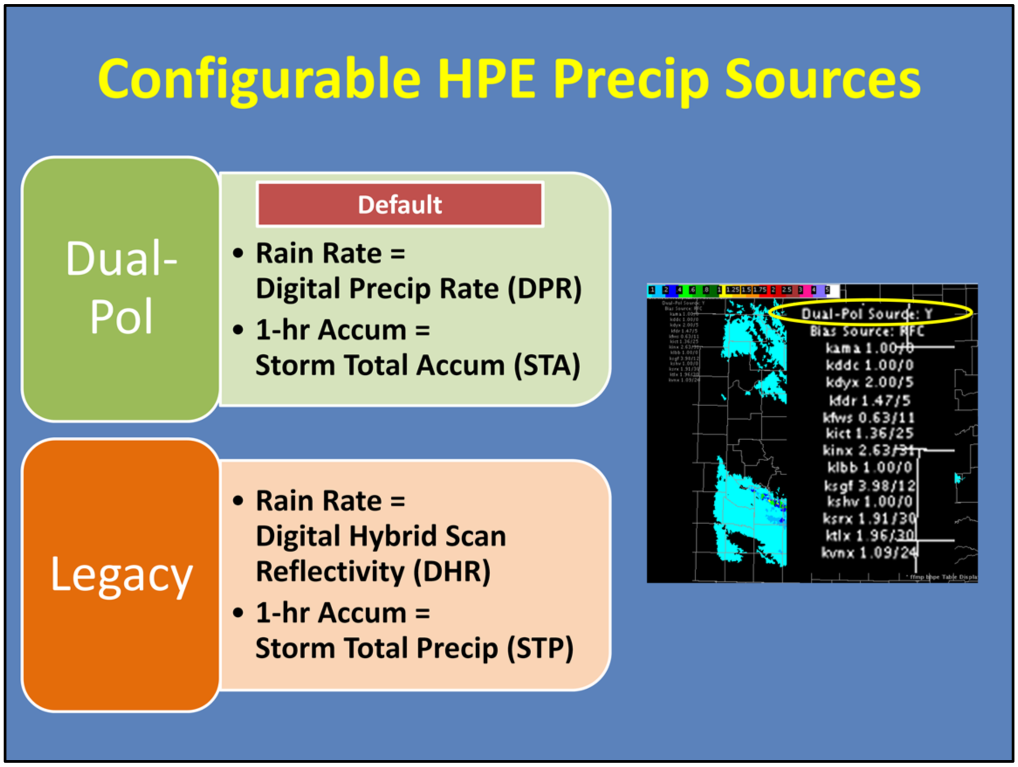

--Dual-Pol or Legacy can be configured by your hydro focal point to create the HPE and Bias HPE mosaics, with Dual-Pol being the default.
--For either precip source, the rain rate product is populated using the respective rate product and the 1-hour accumulation product is created using the storm total product.
--The image below shows an example of the HPE legend in FFMP. And you can see which precip source is being used. Here, it is Dual-Pol. If the source were changed to one of the Legacy products, the letter “N” would be displayed instead.

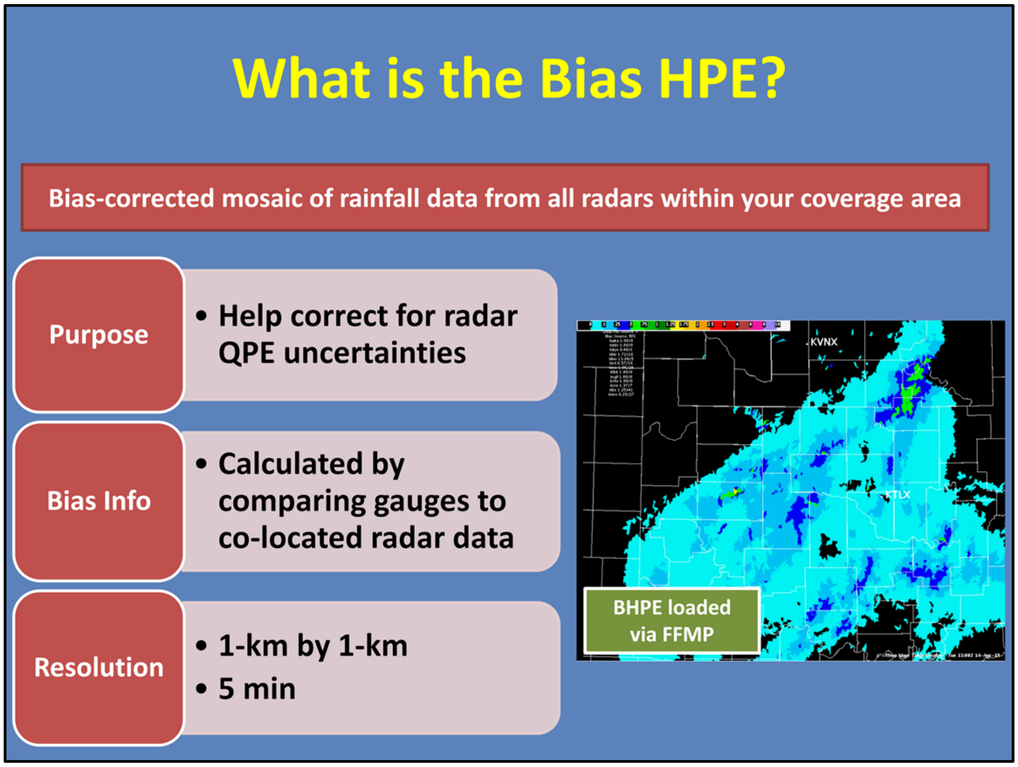

--Bias HPE is the same as HPE, except that a bias correction is made to the precip estimates, meant to help correct for radar uncertainties.
--In short, the bias is calculated by comparing rain gauge information with co-located radar data. This process is detailed in the lesson "Interpreting QPE Bias Information in AWIPS".
--Both products have the same resolution.

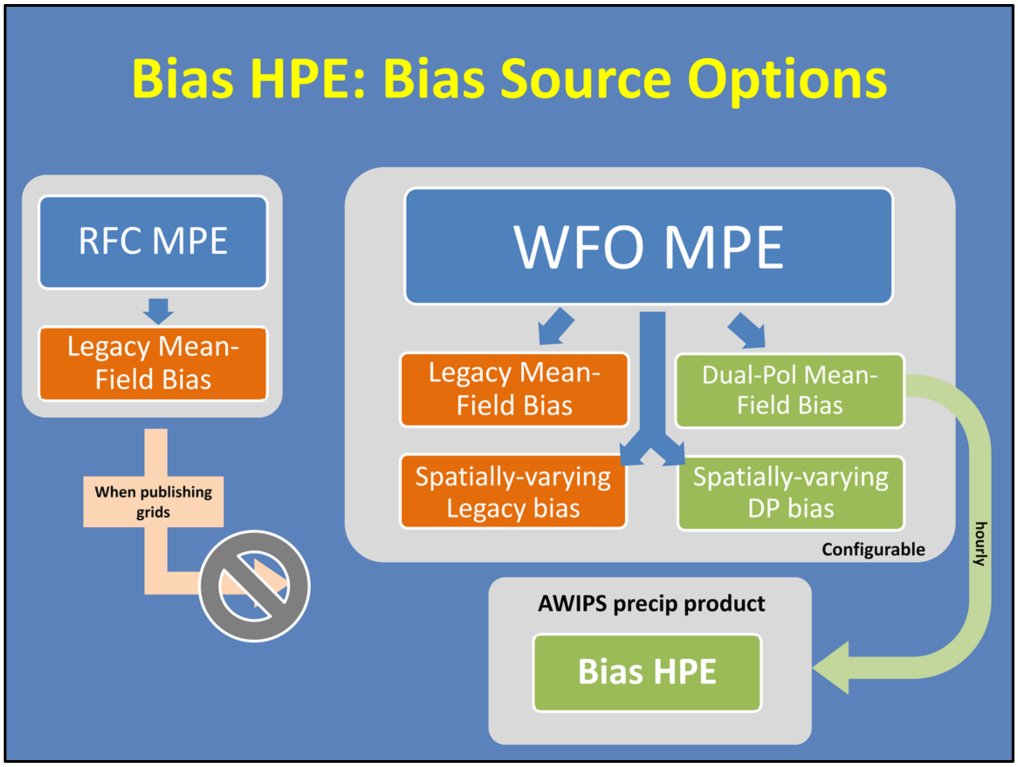

So where is the bias information computed? MPE, or Multi-sensor Precip Estimator, is an AWIPS application used to generate and QC various precip estimates, including gauge-corrected estimates.

Your WFO MPE generates both Legacy and Dual-Pol mean-field biases for each radar, as well as spatially-varying bias grids of each. The mean-field bias options apply one bias factor across each radar's effective coverage area. While the spatially-varying options apply different bias values across the domain based on nearby gauge observations.

Every hour, these biases are sent from the MPE to be used to create the Bias HPE product, amongst other things. Right now, the default is for the Dual-Pol mean-field bias to be used within Bias HPE, but this is configurable. See the VLab reference documentation for more details on this.

The RFCs also have an MPE that creates Legacy mean-field biases that are transmitted to WFOs based on when they publish their QPE grids. However, some RFCs have not been maintaining the code required to transmit the biases. Regardless, since Dual-Pol is now the default for Bias HPE, this doesn't impact the product. It primarily impacts which Legacy mean-field bias factors are displayed in the Legacy precip products.

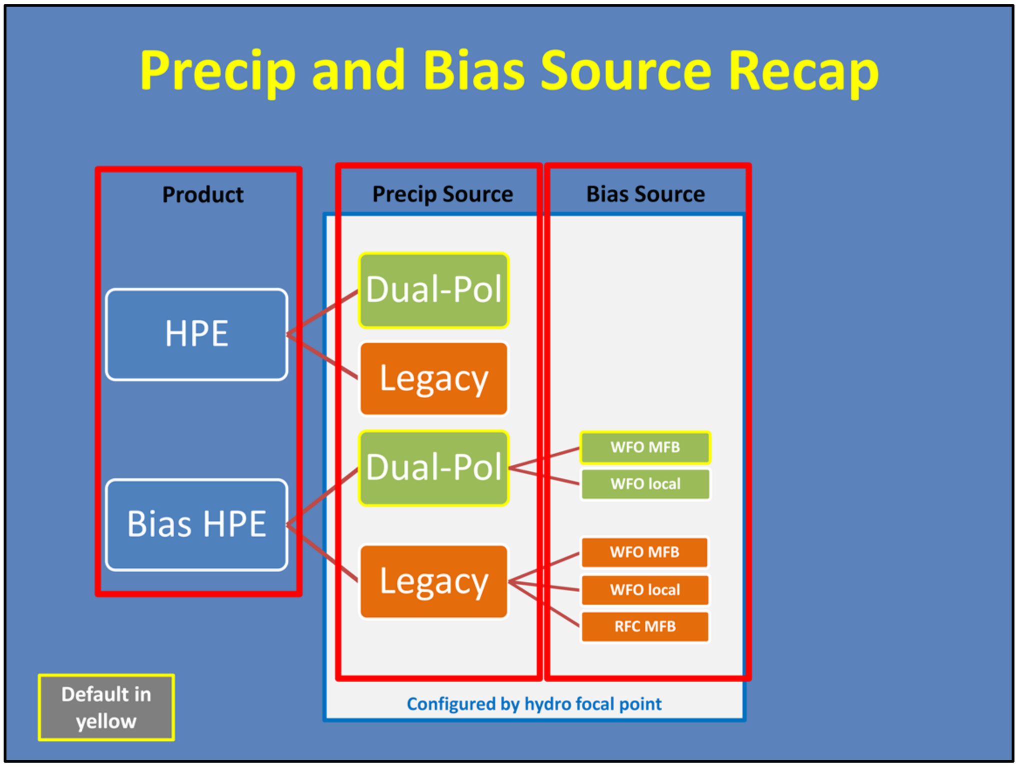

It can be a little confusing to keep track of all of the options between precip sources and bias sources. So here is a little flow chart to help summarize what is available.
--So you, as the forecaster, will have the option of choosing between HPE or Bias HPE to load into FFMP, depending on if you want a bias applied to the mosaic field. From there, the precip source will either be Dual-Pol or Legacy. Dual-Pol is the default, and likely what's configured at most offices. If you choose Bias HPE, then one of two bias sources will be applied to the mosaic: either the WFO mean-field bias or the WFO local bias. Here, WFO MFB is the default.
--All of these choices are configurable by your local hydro focal point, with instructions on the VLab.

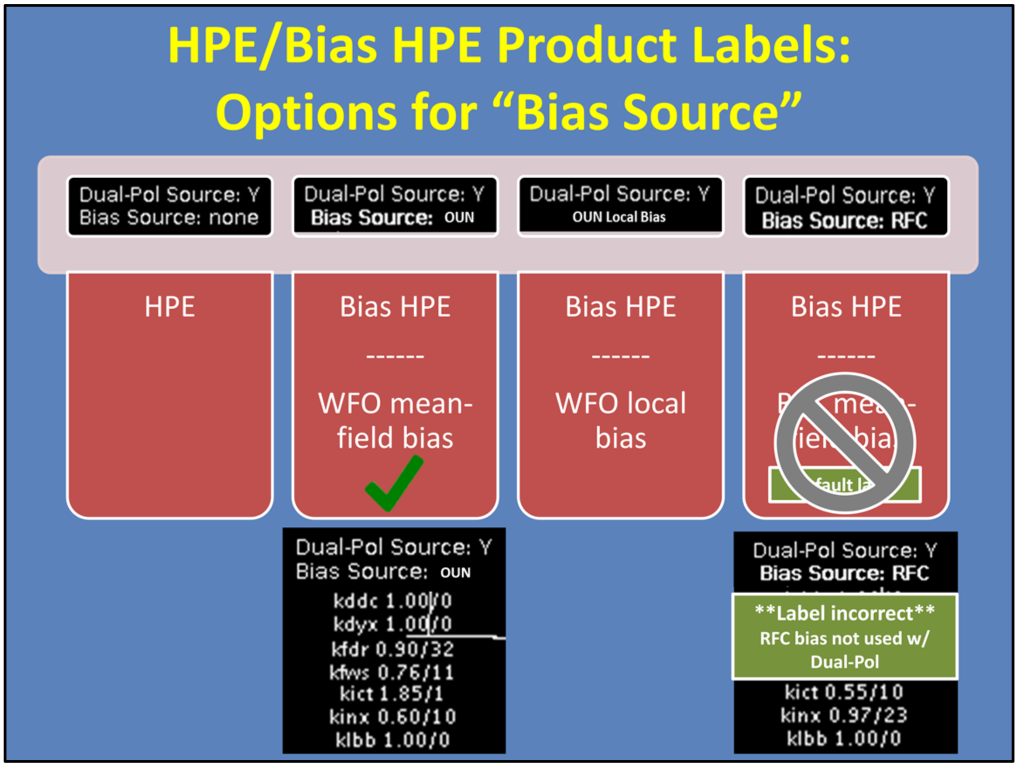

Okay, so let's talk about how this bias information is displayed in the HPE and Bias HPE product labels.

First off, since HPE doesn't have any biases applied, its Bias Source label simply says "none". Easy enough.

Now for Bias HPE, the label changes based on which bias source is configured.
--If the bias source is the WFO MFB, then the three-letter identifier is given (in this example, OUN), and the mean-field bias info is listed for each radar.
--Next, if the bias source is the WFO's spatially-varying bias, then it will say the WFO's identifier, followed by "Local Bias" and nothing will appear underneath. If you recall from the earlier slide, this is because this option does not offer a uniform bias for each radar, so there's nothing to list out.
--Finally, there's the RFC MFB. This is actually the default bias source *label* for Bias HPE. However this is misleading because the RFC bias is not used with the Dual Pol precip source. We have a workaround to fix this on our VLab reference page.

To summarize, Dual Pol is the default precip source for all products, and is labeled correctly. For Bias HPE, the WFO MFB is *generally* used for the bias source, even though the default label may incorrectly say it's the RFC MFB.

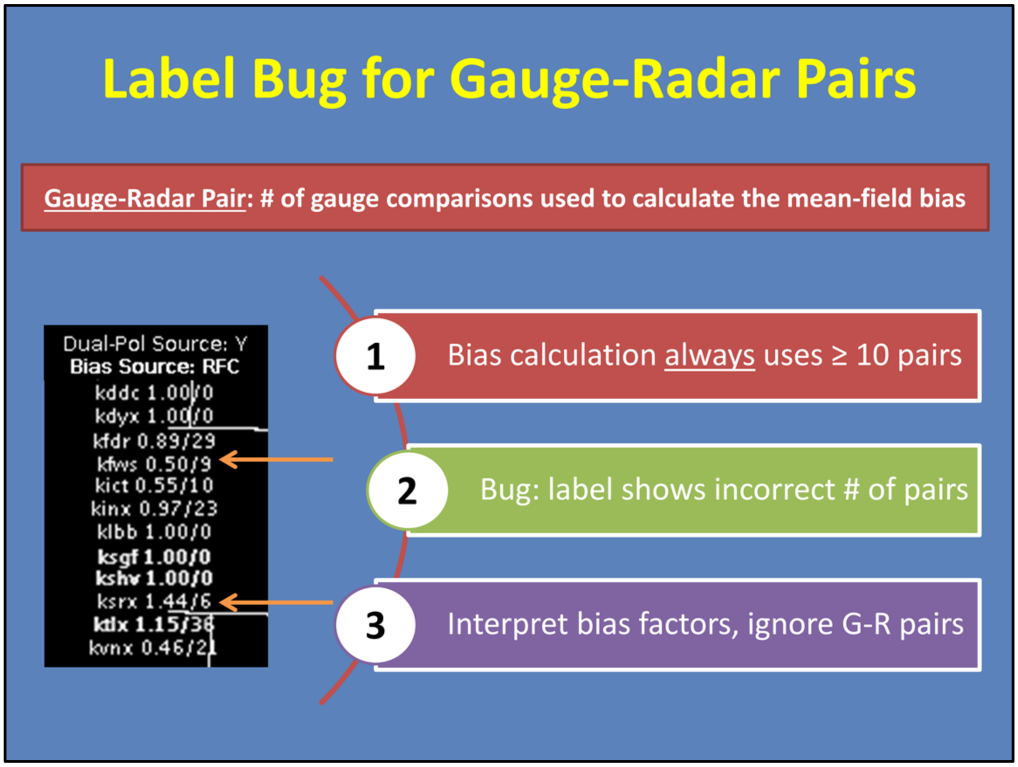

You will learn more in the lesson "Interpreting QPE Bias Information in AWIPS", but briefly, the number next to the mean-field bias factor is the number of gauge-radar pairs used to calculate each radar's mean-field bias. The more gauge comparisons used, theoretically, the more reliable the bias factor should be.

Mean-field biases are *always* calculated using at least 10 pairs. However, there is currently a bug where the label will sometimes say there's less than 10 pairs, as shown here.

The key takeaway is simple...you can interpret your bias factors knowing that they were calculated with the correct number of pairs. Simply ignore the gauge-radar pair number as it can sometimes be wrong.

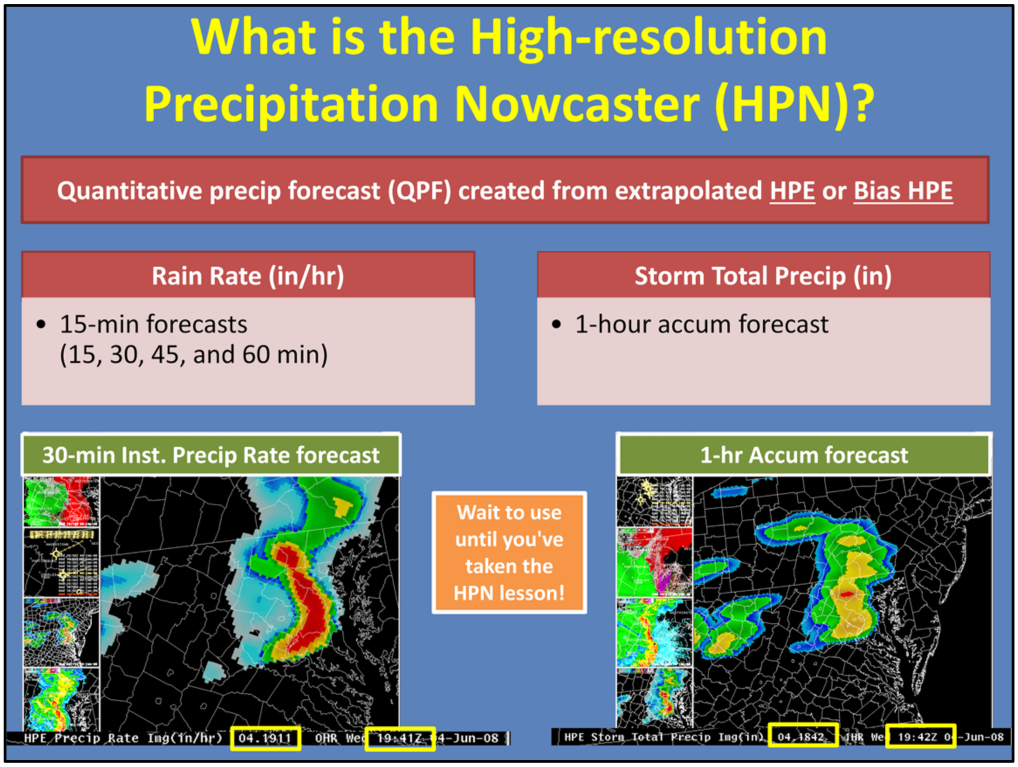

HPE also has a QPF component called the High-Resolution Precipitation Nowcaster (or HPN). This product uses its own feature tracking to extrapolate precipitation forecasts from an HPE or Bias HPE rain rate input. Therefore, the Dual-Pol default carries over into these HPN forecast products, as well.

--The output is QPFs up to one hour in the future. For the rain rate products, four 15-min forecasts are created (at 15, 30, 45, and 60 minutes out). For the accumulated product, the rain rate products can be summed to create a single one-hour forecast.

--HPN products have a slightly lower spatial resolution, at 4km.

--Below is an example of each type of output from the Volume Browser. On the left, is a 30-min forecast of instantaneous precipitation rate for HPE. You can see the current time is 1911 and the product is valid at 1941. On the right is a one-hour accumulation, with the current time being 1842 and the forecast time as 1942.

--Both of these examples are forecasts of HPE, but remember, you can also look at forecasts of Bias HPE, as well.

--Due to the advanced nature of how this product is created, we recommend you wait to use it until you have taken the HPN lesson.

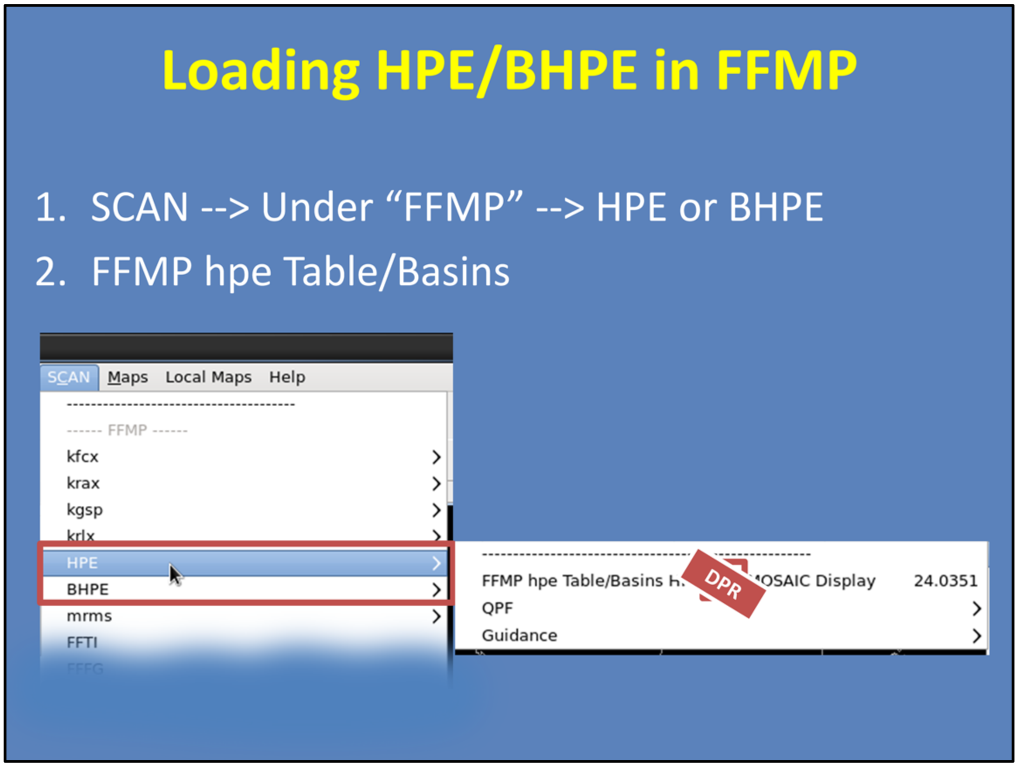

--HPE was designed to make it easier to view multiple radars in one mosaic. This is especially useful in FFMP because you can load just one HPE mosaic, instead of multiple FFMPs for different radars.

--As a reminder, to load HPE via FFMP, go to the SCAN menu and click the HPE (or Bias HPE) sub-menu. From there, open the FFMP Table/Basins display.

--Remember that the default HPE source is Dual-Pol, though the label incorrectly says "DHR". To be correct, it should read "DPR".

--If you are interested in how to load these products, as well as the HPN grids, via the Volume Browser, please refer to the VLab reference material.

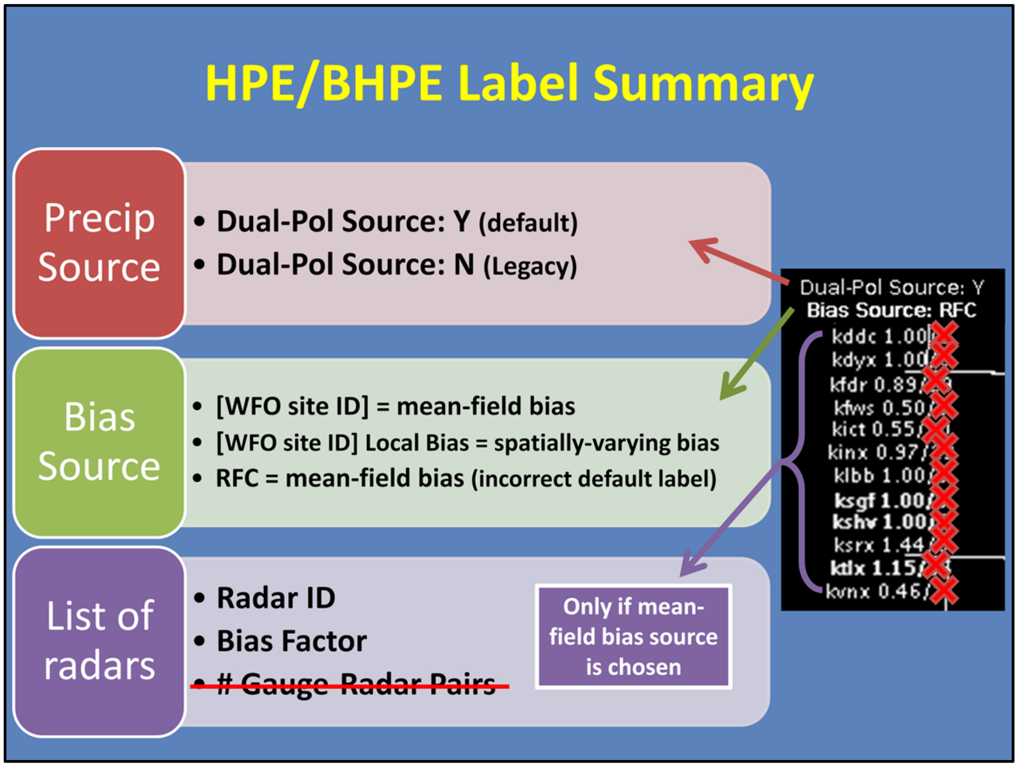

--Here is a summary of how to read the label information for HPE and Bias HPE.
--The first line specifies the precipitation source. If it says "Dual-Pol: Y", then Dual-Pol is being used. This is the default. If it is an "N", then Legacy is used.
--The second line specifies the bias source. There are three options here: If it is the WFO's mean-field bias, then it will state the WFO's three-letter identifier. This is the actual default. If it is the local spatially-varying bias, then it will specify the WFO's three-letter identifier with "Local Bias". If the RFC mean-field bias is used, then it is says "RFC". Currently, the default label incorrectly says "RFC" when it is generally using the WFO mean-field bias.
--Finally, if one of the mean-field bias sources is configured, then the bias information for each radar is listed.

Remember, that the gauge-radar pair info can be labeled incorrectly, so simply ignore these values.

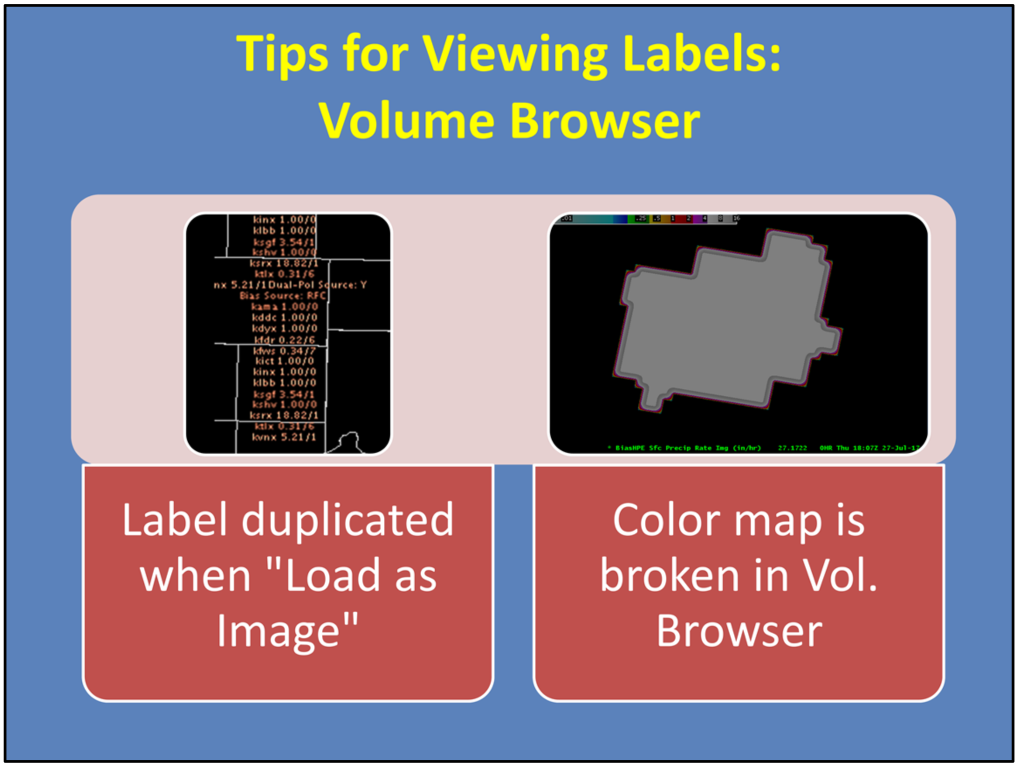

--There are a few things to be aware of when viewing HPE and Bias HPE.
--When viewing in the Volume Browser, the label gets duplicated when you “Load as Image”. This can make it a little difficult to read, but none of the text is overlapping, so the important information is still visible.
--Also note that the color map is broken in the Volume Browser display for most localizations (for both HPE and Bias HPE). But we have a fix for this problem on the VLab.

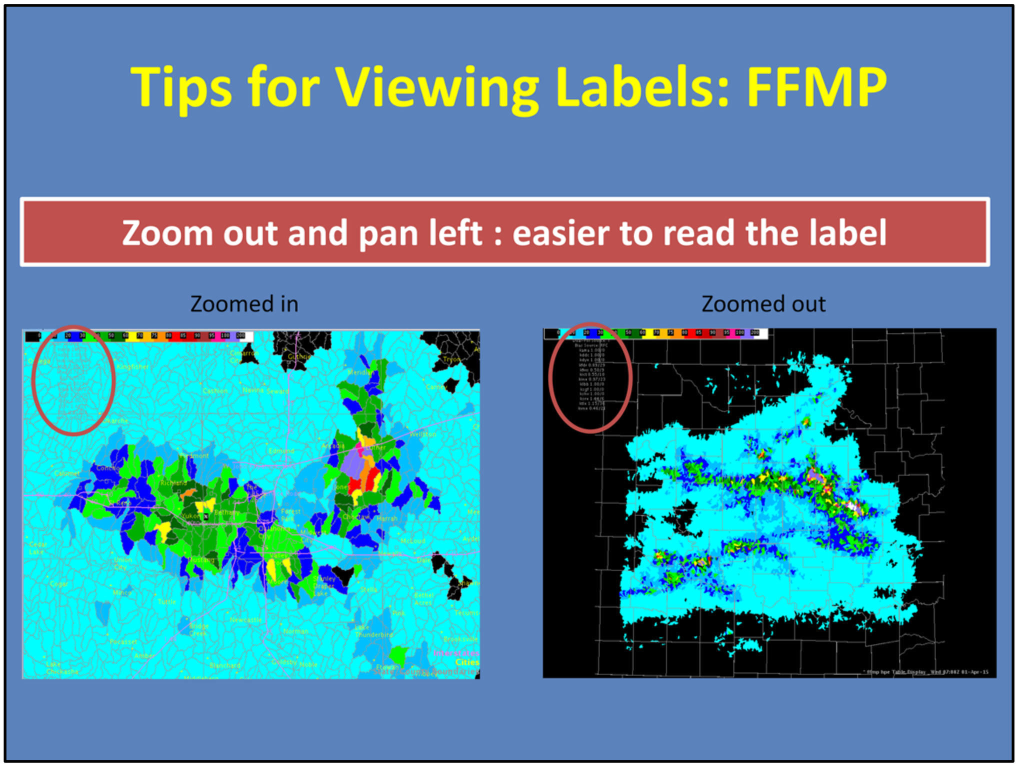

--When viewing in FFMP, make sure to zoom out and pan to the left border in order to better read the label. When zoomed in, the label may be obscured by the basin fill. But if you zoom out, you can move it to a blank area to read easier.

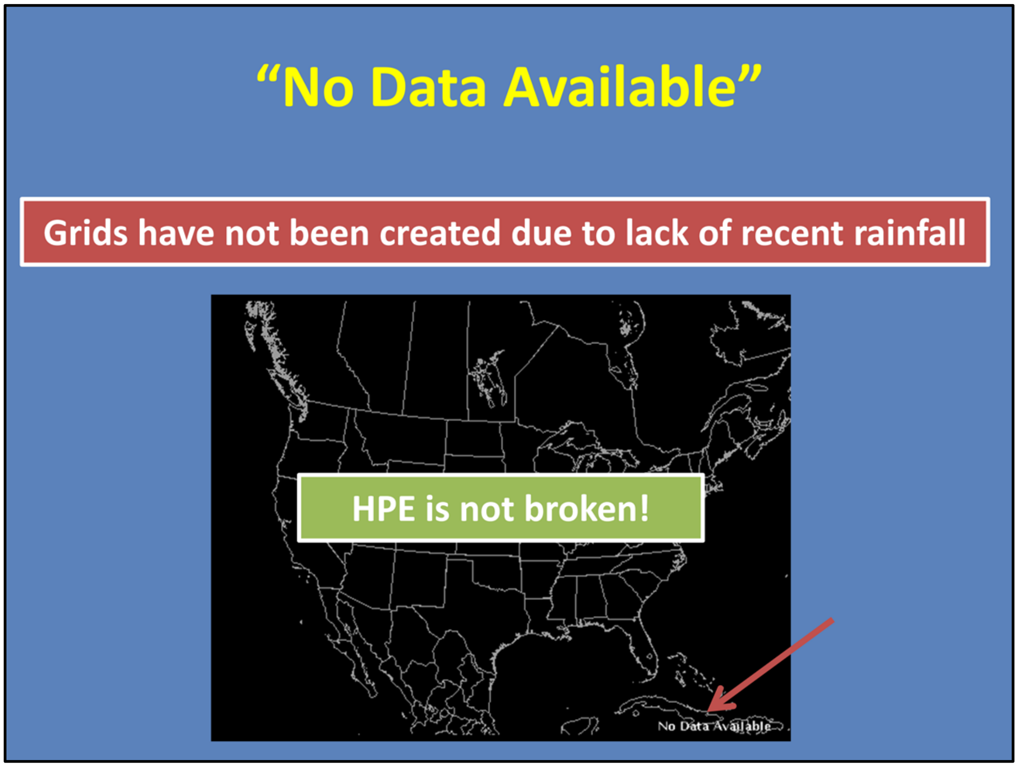

If you try to open HPE or Bias HPE when there has been no rainfall in your domain, you will get a blank pane with "No Data Available" in the bottom-right. This does not mean that HPE is broken! It simply means that the grids have not been created due to the fact that there has not been any recent rainfall. Just wait until the next rainfall occurs in the domain, and try opening again.

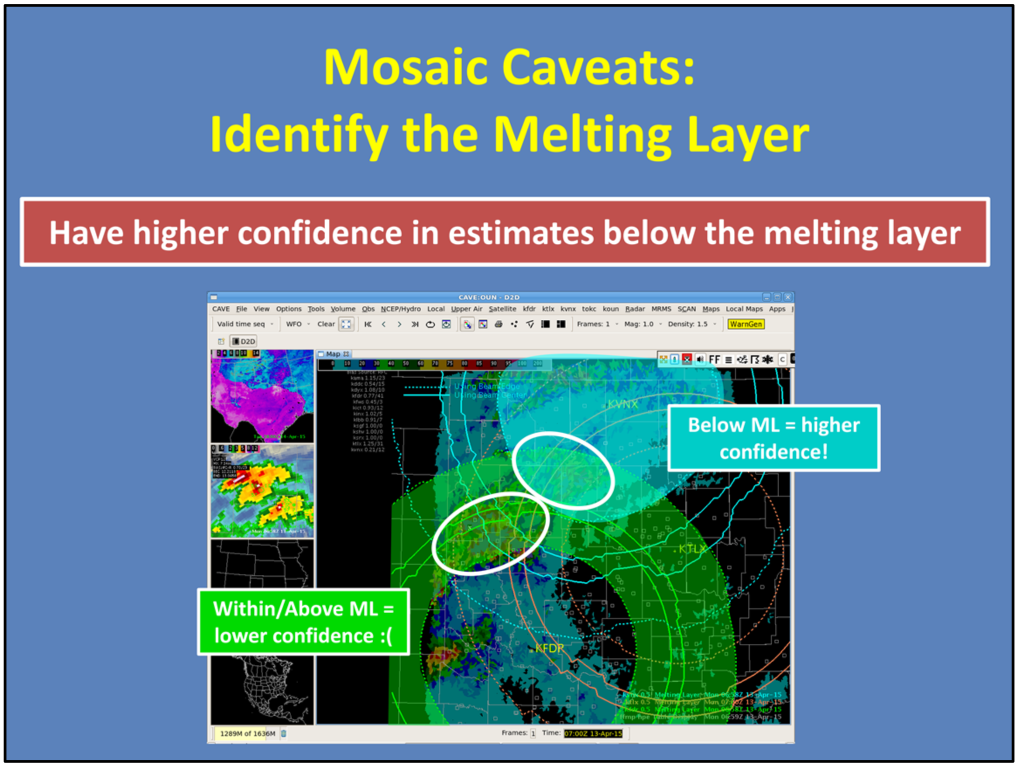

There are some caveats to using any Dual-Pol mosaic, like HPE and Bias HPE.

--For one, identify the melting layer by using the low-level melting layer overlay. Have higher confidence in estimates below the melting layer. Within and above the melting layer, your estimates are going to change, and they won't be as good. This is due to the HCA assigning mixed phase or ice classifications at these heights, and thus, the Z-R relationships are being matched to those precip types, even if they are liquid when hitting the ground.

--Here is an image of HPE. Let's consider the circled areas that may be experiencing flash flooding. By overlaying the melting layers for the three nearby radars, you can see the northeastern circled area is closest to the most northern radar, or the one with the blue ML circles. These values are located below the ML, so you can have higher confidence in them. However, the area to the southwest is within the ML circles for both radars. Thus, this area has more uncertainty.

Back to Top

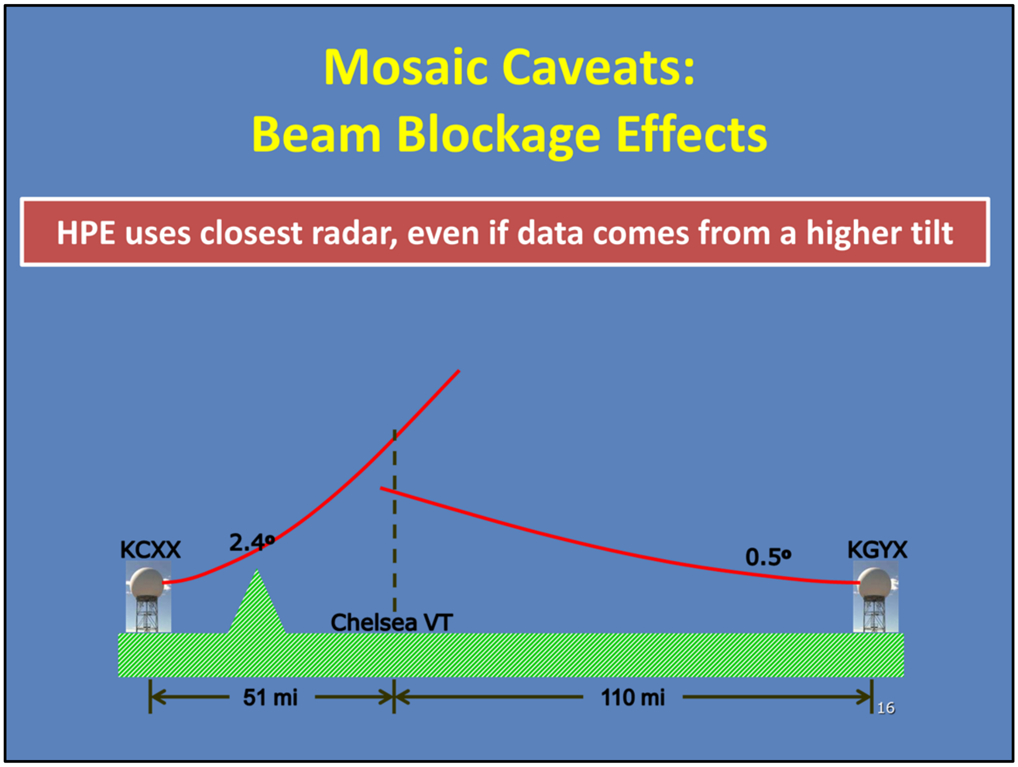

Secondly, watch for artifacts in areas of beam blockage. In this example, HPE only knows that KCXX is the closer of the two radars, so it uses the KCXX precip estimates at the given point. However, we see its precip estimates are coming from the 2.4 deg tilt. KGYX actually has the lower altitude estimate at this grid point. So how do we fix this?

Well, you can force HPE to use the farther radar with what's called a misbin file. For more information on this topic, please see the Resources tab for a presentation from Greg Hanson.

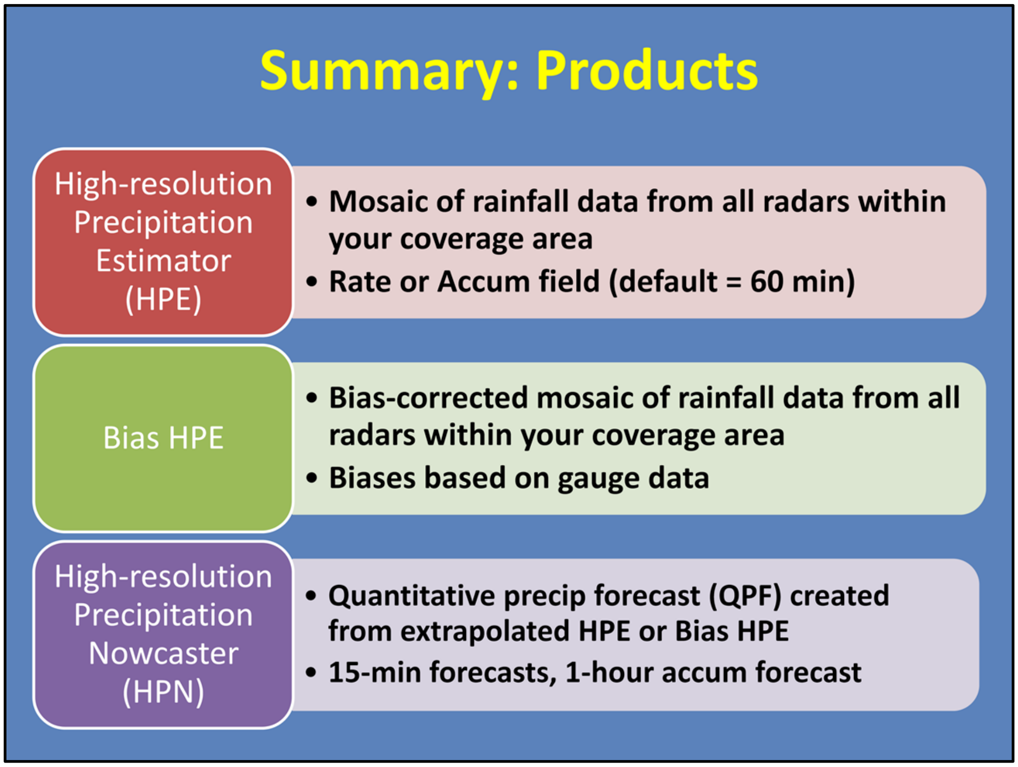

So let's quickly summarize the products we discussed in this lesson.

1. The High Resolution Precipitation Estimator (or HPE) is a gridded rainfall mosaic, based on single radar sources. It can be created as an instantaneous rain rate or accumulation field, with the default accumulation being one hour.
2. The next product is the Bias HPE. This is similar to HPE, except that biases have been applied to the radar estimates to correct for uncertainties. These biases are calculated using available gauge data.
3. Finally, there is the High Resolution Precipitation Nowcaster (or HPN). This product is an extrapolated QPF, based on either the HPE or Bias HPE product. The forecast is available for 15-minute intervals up to one hour, or for a single one hour accumulation.

All three of these products update every 5 minutes.

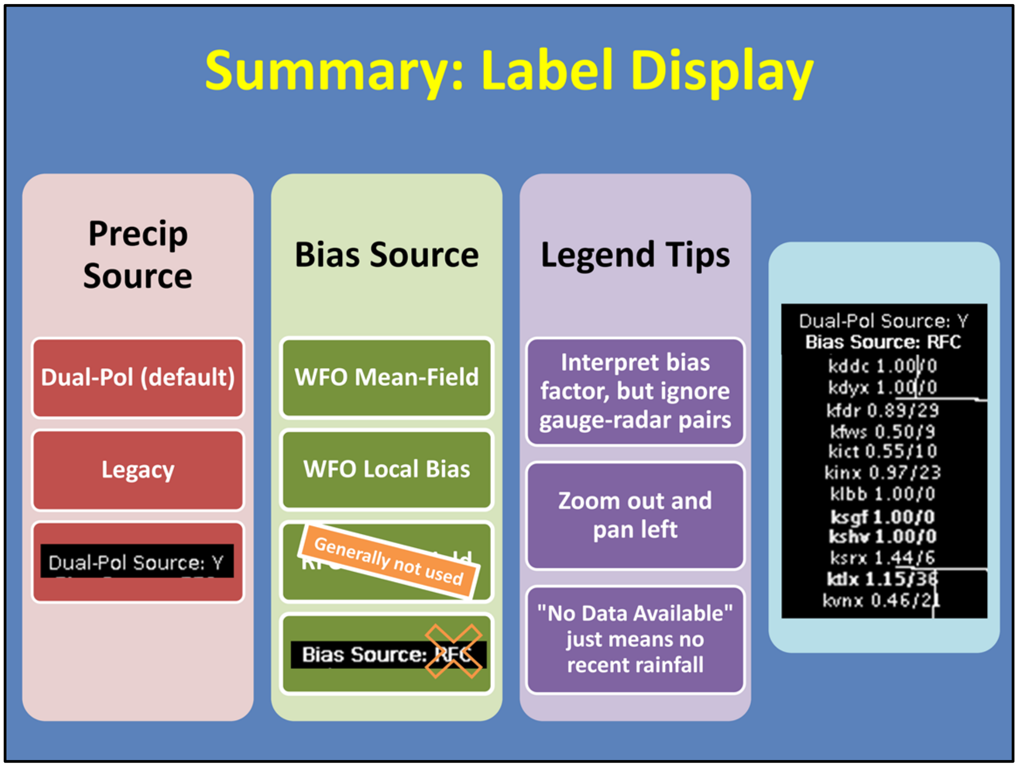

The labels for HPE and Bias HPE can be used to interpret valuable information about the products.
--For one, you can see which precip source is being used, either Dual-Pol or Legacy, with Dual-Pol as the default.
--For Bias HPE, you can also determine the bias source. The default is the WFO mean-field bias, but you can also use your WFO local bias. Keep in mind, if your label says RFC (and you're using Dual-Pol), it's incorrect because the RFCs do not send Dual-Pol mean-field biases.
--Remember that you can interpret the bias factors from the legend, but ignore the gauge-radar pair information.
--In FFMP, you can view the legends easier by zooming out and panning to the edge.
--Also, if you see a blank pane and “No Data Available” when loading in FFMP, it means there hasn’t been enough recent precip for the HPE grids to be calculated. Simply try loading again once rainfall has occurred in the area.

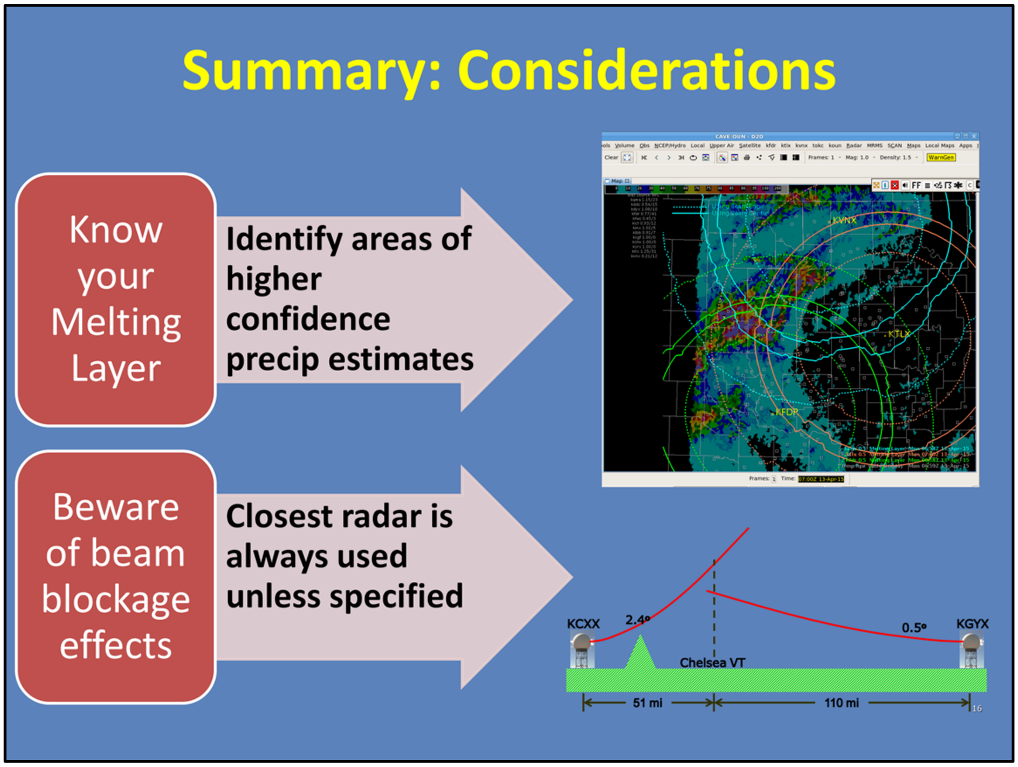

Finally, there are some important considerations to better apply HPE and Bias HPE.
--First, you should overlay the low-altitude melting layer algorithm to identify areas of higher confidence precip estimates.
--Also, beware of beam blockage in complex terrain, as it may affect the QPE estimates.

This is the end of this lesson. When you are ready, please move onto the next slide to take the quiz and receive credit on the LMS.

Back to Top

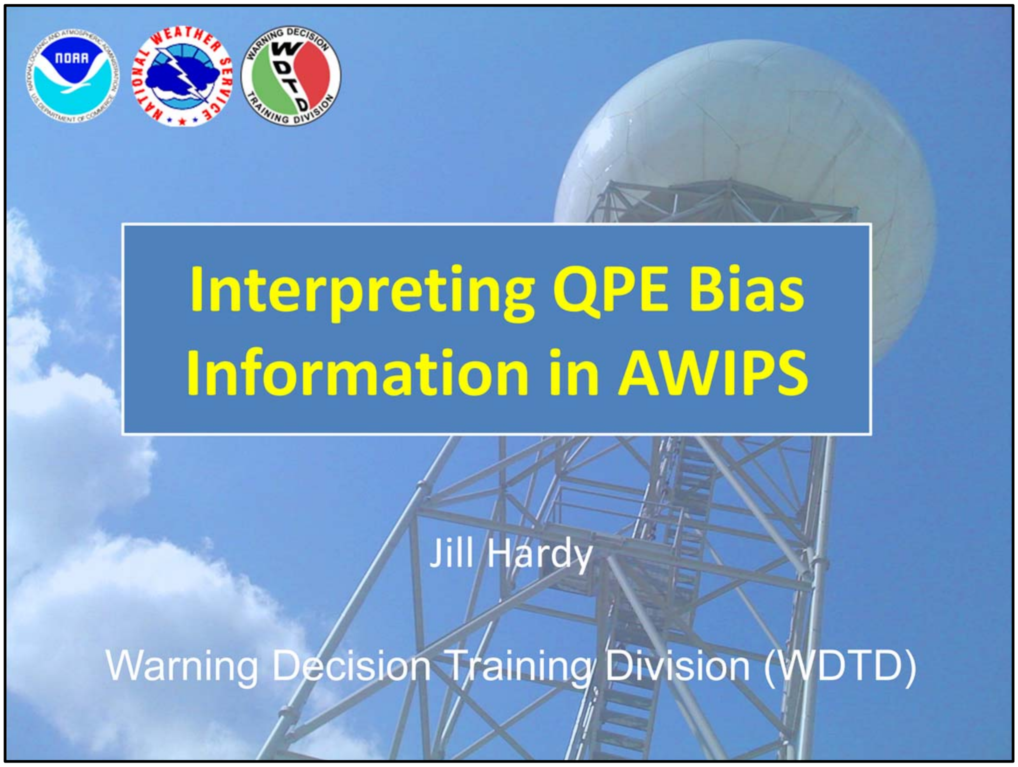

Hi, my name is Jill Hardy and welcome to this lesson on how to interpret QPE bias information in AWIPS.

Back to Top

Learning Objectives

- By the end of this lesson, you will be able to:
  - Determine the bias option(s) for each precip source
  - Define a mean-field bias
  - Interpret a mean-field bias from a product legend
  - Identify times when you should be cautious of interpreting mean-field bias information
  - Identify ways to view mean-field biases in AWIPS
  - Identify when mean-field biases are applied
  - Interpret MRMS bias information
  - Interpret a manual bias calculation with a gauge
  - Identify best practices when performing a manual bias calculation

Here are the learning objectives for this lesson. We will cover everything you need to know about what the biases are, how to interpret them, and a little on the application.

Default Available AWIPS Precip Sources

More on QPE sources: "Choosing Your Precip Sources"

QPE Source	Z-Rs	Bias Option
Legacy	**Single Z-R** (set at RPG)	**Mean-field**
		Spatially-varying (not commonly used)
Dual-Pol	**Spatially-varying** (based on precip type)	**Mean-field**
		Spatially-varying (available in Bias HPE; weighted based on distance from gauges)
MRMS	**Spatially-varying** (based on precip type)	**Spatially-varying** (weighted based on distance from gauges)

Before we dig into the guts of interpreting biases, let's talk about the three available precip sources in AWIPS for which biases can be computed.

These precip sources are Legacy, Dual-Pol, and MRMS. To learn more on how to determine which source is best in any given situation, please see the "Choosing Your Precip Sources" lesson.

But for the purposes of this lesson, we're going to focus on how the default bias options are computed for each one.
--Legacy QPEs are computed using one Z-R relationship across the whole radar umbrella, set at the RPG. Its default bias option is what's called a "mean-field bias". As the name implies, this is one averaged value that can be applied across the radar field.
--Moving onto Dual-Pol, it's calculated using spatially-varying Z-R relationships based on precip type at each pixel. The mean-field bias is an option, but having a single number for a whole radar umbrella doesn't allow for bias differences in different precip types. The spatially-varying bias allows for biases to change across the radar umbrella, but it requires a higher density gauge network to be useful.
--Finally, MRMS is a QPE mosaic that also has spatially-varying Z-Rs based on precip type. And it, too, has implemented spatially-varying logic in the bias-corrected product.

Knowing this, the rest of the lesson is going to describe the two bias options, where to find them, and how to interpret them for each precip source. Let's get to it!

Back to Top

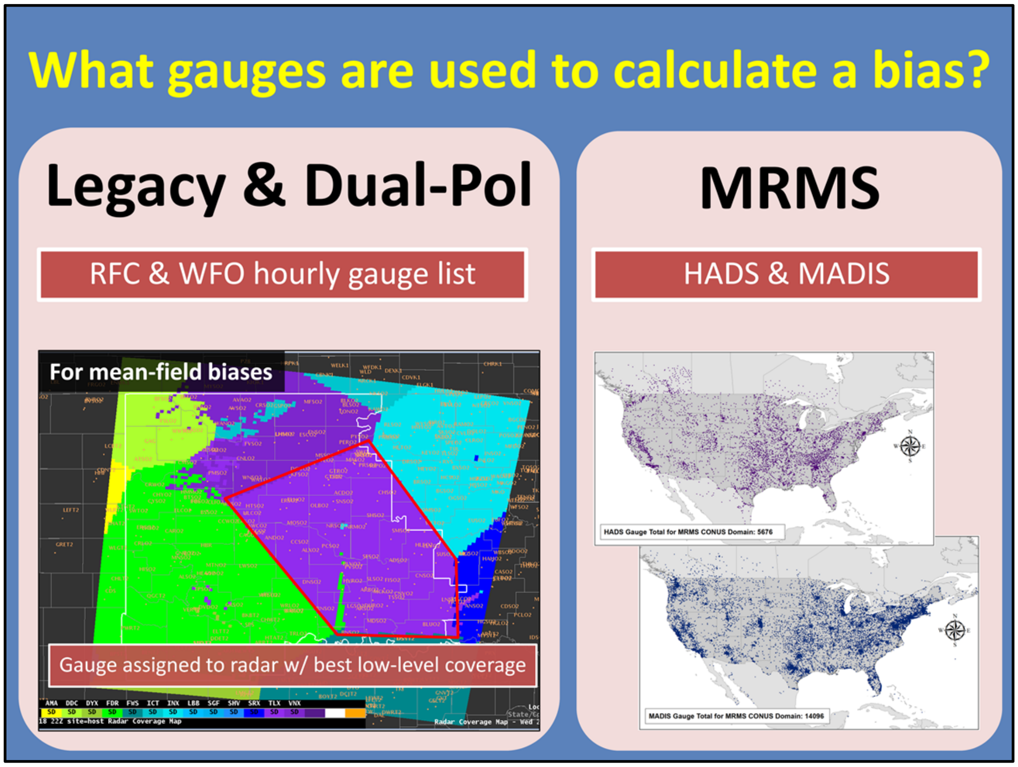

The idea of a bias correction inherently relies on observations that calibrate the radar estimates. So the first question to ask is: "Well, what gauges are even used to calculate a bias?"

For Legacy and Dual-Pol, a list hourly gauges is maintained at both the RFCs and WFOs. RFCs tend to better QC their list, so when possible, their information is used. When calculating each radar's mean-field bias, the hourly gauges are assigned to the radar with the best low-level coverage. Here is an MPE display showing the low-level radar coverage for WFO Norman, with hourly gauges overlaid. So, for example, any gauge that falls within a light purple gridpoint will be used to calculate the mean-field bias for KTLX. To learn more about viewing the gauge info in the MPE perspective, visit the VLab references.

For MRMS, the gauge networks frequently update and improve. Currently, MRMS uses a combination of the HADS and MADIS networks. To learn more, please visit our MRMS Hydro Products course.

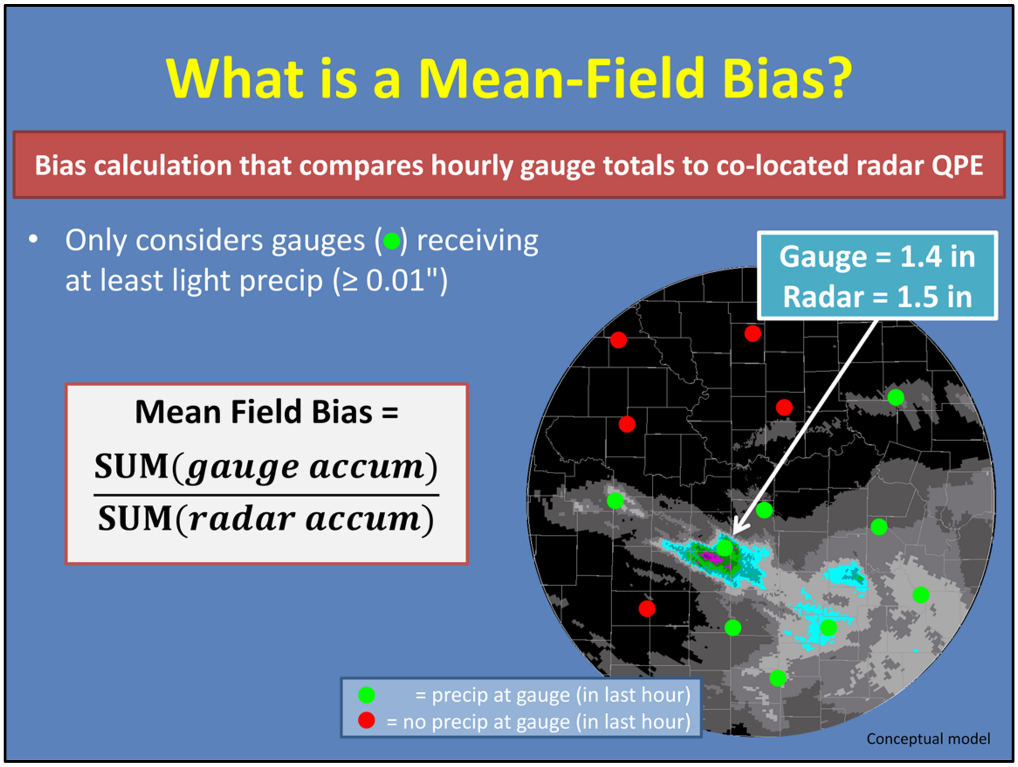

Okay, let's start off by defining what a mean-field bias is. It is a bias calculation that compares hourly gauge totals to co-located radar QPE. Its purpose is to help the forecaster interpret whether the radar estimates are running higher or lower compared to gauge observations.

Consider this radar umbrella, showing Legacy one-hour precip totals. Remember that this is NOT the coverage map used in the calculation, that's the MPE grid shown on the previous slide. But for the sake of simplicity, I will continue using radar examples like these throughout the lesson. So getting back...Each dot represents a gauge location, with green gauges getting precip in the last hour and red gauges not. The mean-field bias calculation only considers gauges receiving at least light precip, so only the green dots.

Let's focus in on one gauge. The one-hour total at this gauge was 1.4 inches. The co-located radar QPE value was 1.5 inches during the same accumulation period. This is called a gauge-radar pair. And at this location, the Legacy QPE was slightly higher than the observation.

So the mean-field bias calculation is this: It takes the sum of all the gauge accumulations in the hour (so the totals at each green dot), and divides it by the sum of all the radar QPE accumulations at those locations. We'll discuss what that means on the next slide.

Back to Top

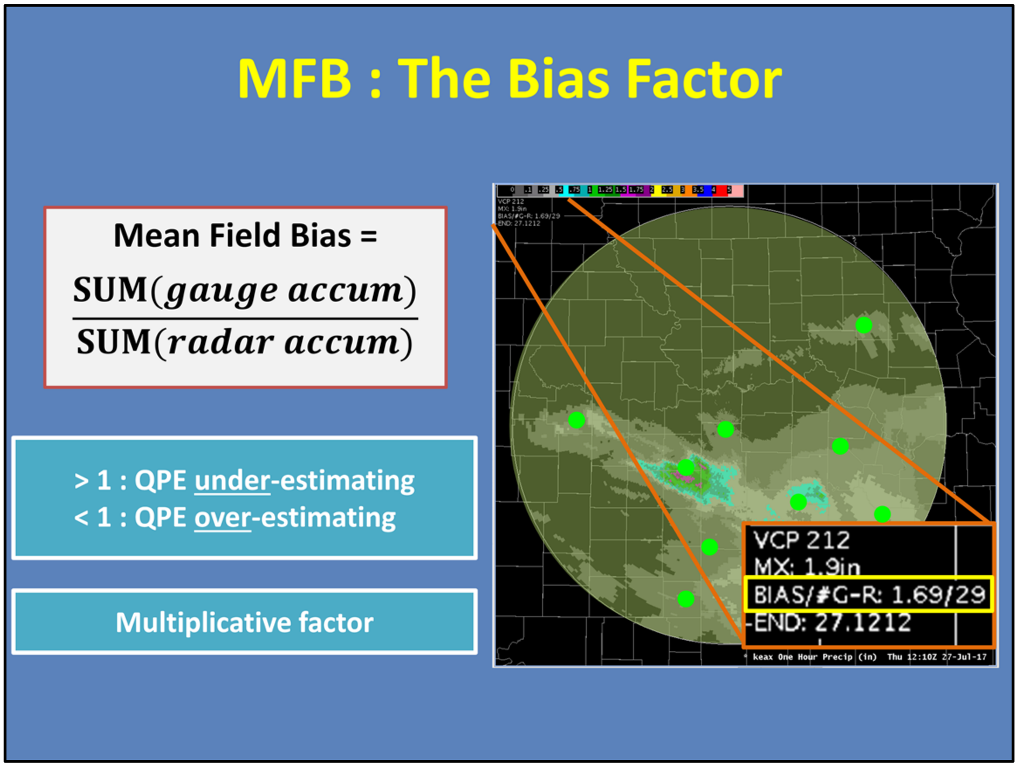

When viewing mean-field bias information in AWIPS, it will be displayed in one of a few ways. The first is in the upper-left corner of your Legacy one-hour and storm-total precip products, as shown in this image. We will talk about the other ones later.

There are always two pieces of information given in the bias display: the bias factor itself, and the number of gauge-radar pairs. First, let's discuss the bias factor.

The bias factor is the result of the calculation shown on the previous slide, and duplicated here. If all the gauges together have a higher total than radar estimates at those locations, the bias will be greater than 1, and the precip source is under-estimating. If the gauges have less than radar, then the bias is less than 1, and the precip source is over-estimating. For example, in this image, the bias factor is 1.69. Therefore, the precip source is under-estimating.

It is a multiplicative factor. So even though, in reality, the bias is calculated over a smaller domain than shown here, it is meant to be multiplied across the entire radar umbrella. So here, the algorithm recommends multiplying the QPEs by 1.69 to bring them up to where they should be. By default, the bias factor is not applied at the RPG, just given in these Legacy precip legends for reference.

Back to Top

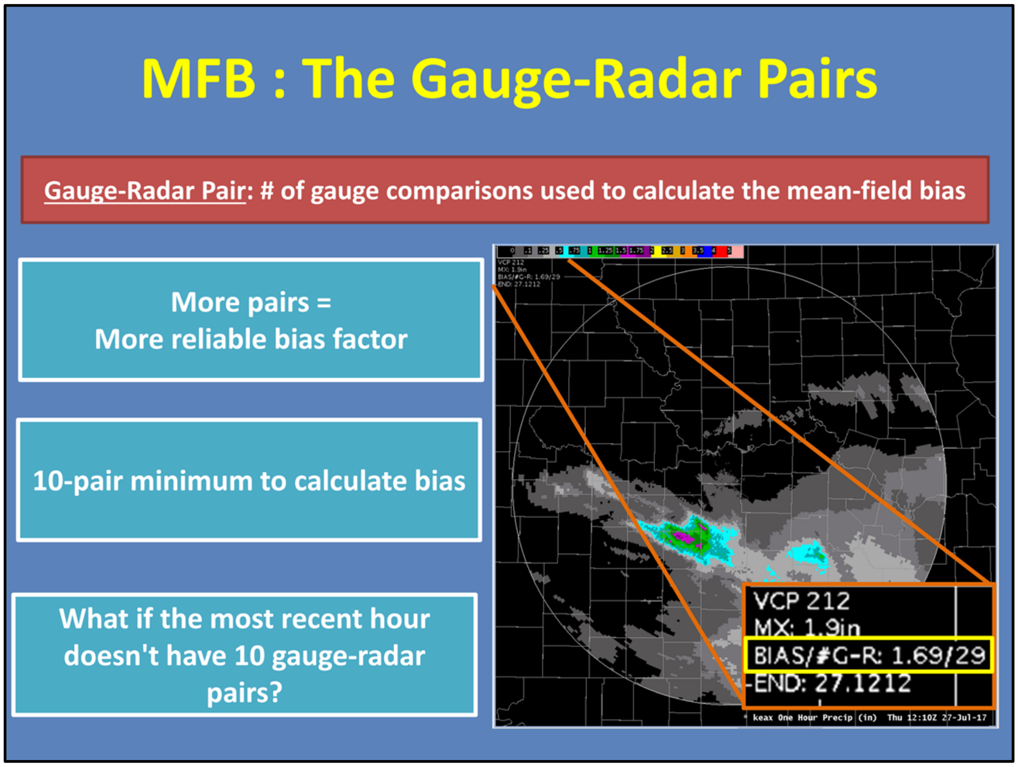

The second component is the number of gauge-radar pairs used to calculate the bias. This is provided in order to add confidence to the bias factor. The more pairs used, theoretically, the more reliable the bias should be. In this example, there were 29 gauge-radar pairs used to get the 1.69 bias factor.

In order to calculate the bias, there is a minimum of 10 gauge-radar pairs needed. But what happens when the most recent one-hour of rainfall doesn't hit enough gauges?

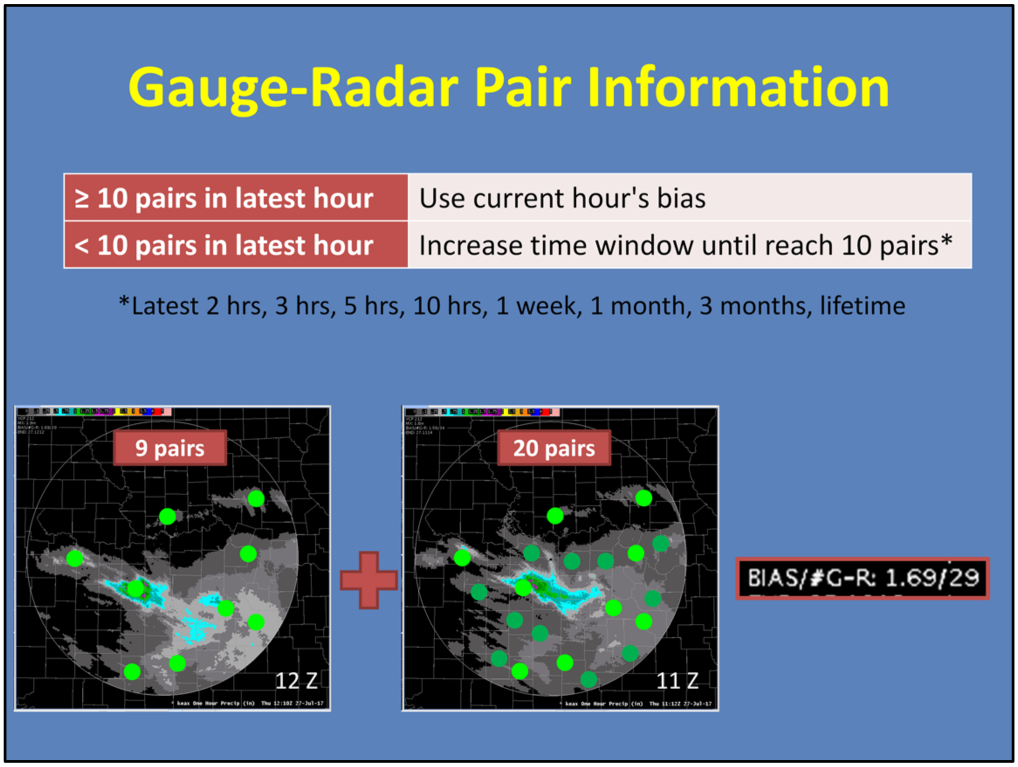

So, if there are 10 or more G/R pairs with accumulations in the latest hour, then the current hour's bias factor is used. But if there are less than 10 G/R pairs in the latest hour, then the algorithm increases the time window to collect more gauge information. If it reaches 10 or more gauges in the latest two hours, it creates a bias factor. If it is not reached by then, it will look at the latest 3 hours, and so on...even potentially very far in the past.

Let's consider our example. In the current hour, there are only 9 gauges that received precip. That's not enough to calculate a bias factor. So the algorithm looks back an hour ago, to 11Z, and finds that 20 gauges hit during that hour, including some of the same ones. It can now use all of these pairs (from both hours!) to create the bias. Hence why there are 29 pairs in the display.

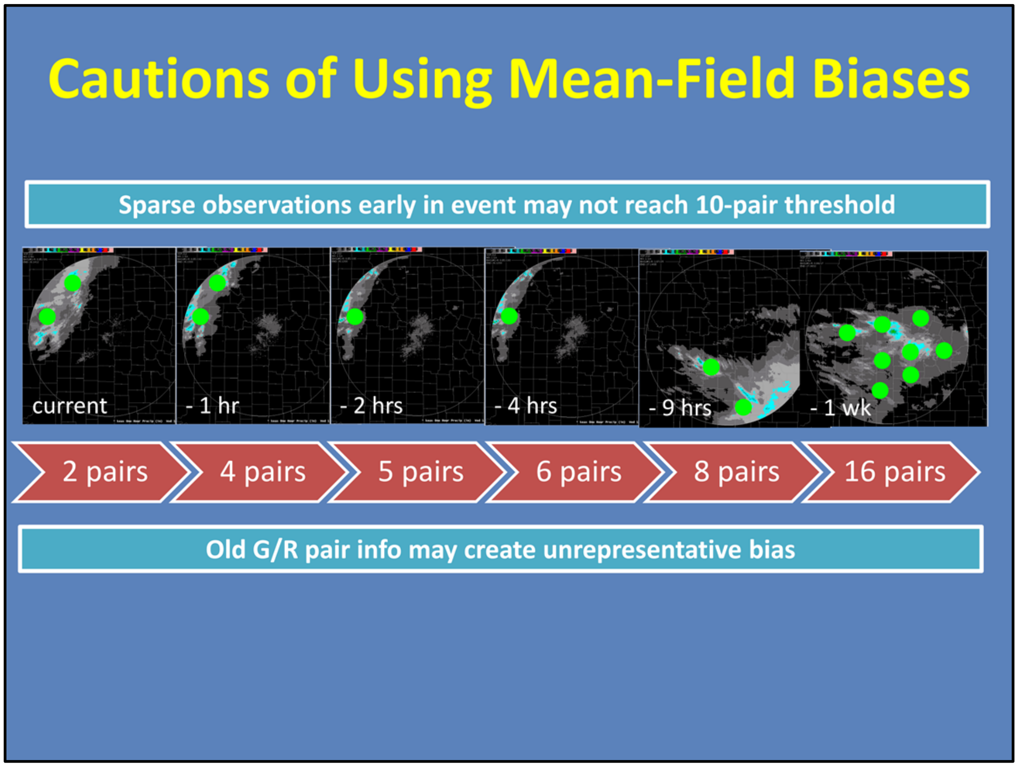

A caution of this method is if you are early in an event, you may not have a lot of recent rainfall to trigger gauge accumulations. Therefore, the algorithm may be looking VERY far back in time to reach the 10 G/R pair threshold. Take, for instance, this example. Since the rainfall is just moving into the radar's umbrella, the most recent several hours don't have enough gauge hits to calculate a bias. There are only 6 total pairs in the latest 5 hours. Only until the algorithm looks back to one week ago, did it get over that 10 pair threshold.

At this point, half of your gauge information is quite old. And old G/R pair info may not be representative of the current event, thus creating an unrepresentative bias.

Back to Top

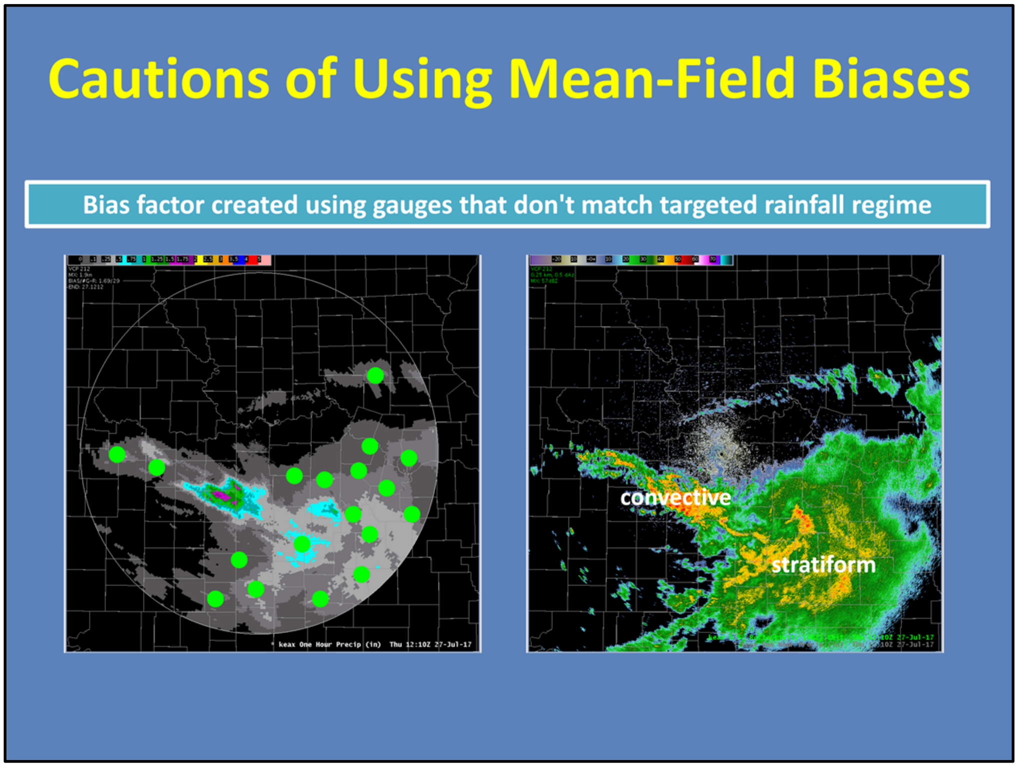

Another caution is that the bias factor could be created using most, if not all, gauges that don't match your targeted rainfall regime.

In our example, it may be more important to properly bias correct the convective area, since it's creating the highest rainfall totals. However, maybe the stratiform region is where all the gauge hits are for that hour. And what if those very light totals create a bias factor that is unrepresentative of the convective region? Would you necessarily want to apply it there? Unfortunately, you don't get that choice. That's where the term "mean field" comes into play.

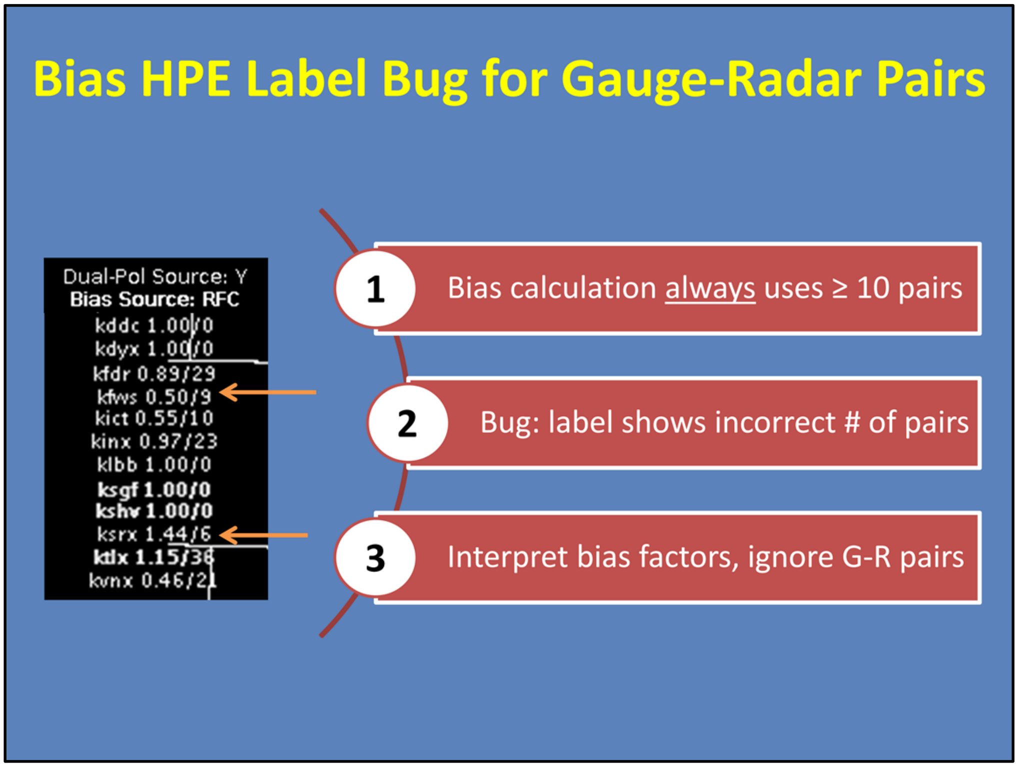

Before I move on, I should mention there's currently a bug in the Bias HPE label related to the gauge-radar pairs.

Mean-field biases are *always* calculated using at least 10 pairs. However, the label will sometimes say there's less than 10 pairs, as shown here. The key takeaway is simple...you can interpret your bias factors knowing that they were calculated with the correct number of pairs. Simply ignore the gauge-radar pair number in Bias HPE as it can sometimes be wrong.

Back to Top

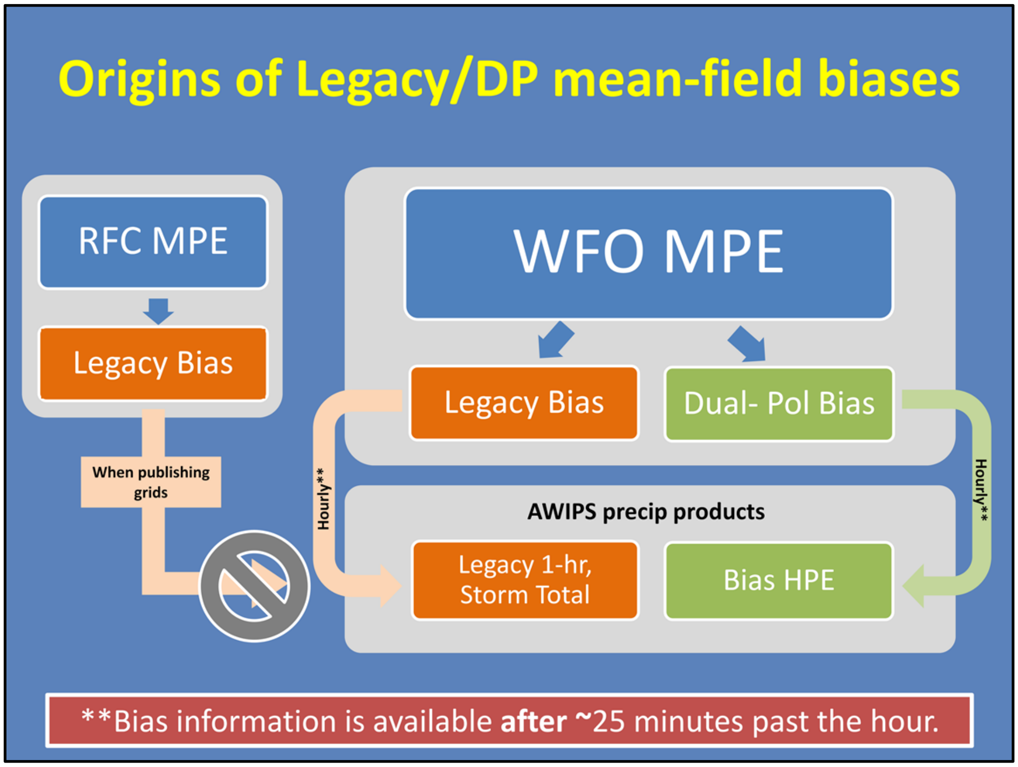

So where is the mean-field bias information computed? MPE, or Multi-sensor Precip Estimator, is an AWIPS application used to generate and QC various precip estimates, including gauge-corrected estimates.

Your WFO MPE generates both Legacy and Dual-Pol bias information, and then disseminates it hourly to your Legacy precip legends, and Dual-Pol Bias HPE. More on this in the next slide.

The RFCs also have an MPE used for their daily tasks, including the creation of just Legacy bias information. These Legacy biases are supposed to be sent when the RFCs publish the grids, but currently, some RFCs do not have MPE configured to send these out.

Theoretically, the RFC biases are thought to be better since RFCs are QC'ing their precip and gauge data constantly. So once the RFCs are configured to send biases to the WFOs, these RFC biases will take precedence over the WFO Legacy biases.

Also, it's important to know that your new bias information will NOT appear until after about 25 minutes past the hour. This is due to known latencies in gauge data collection.

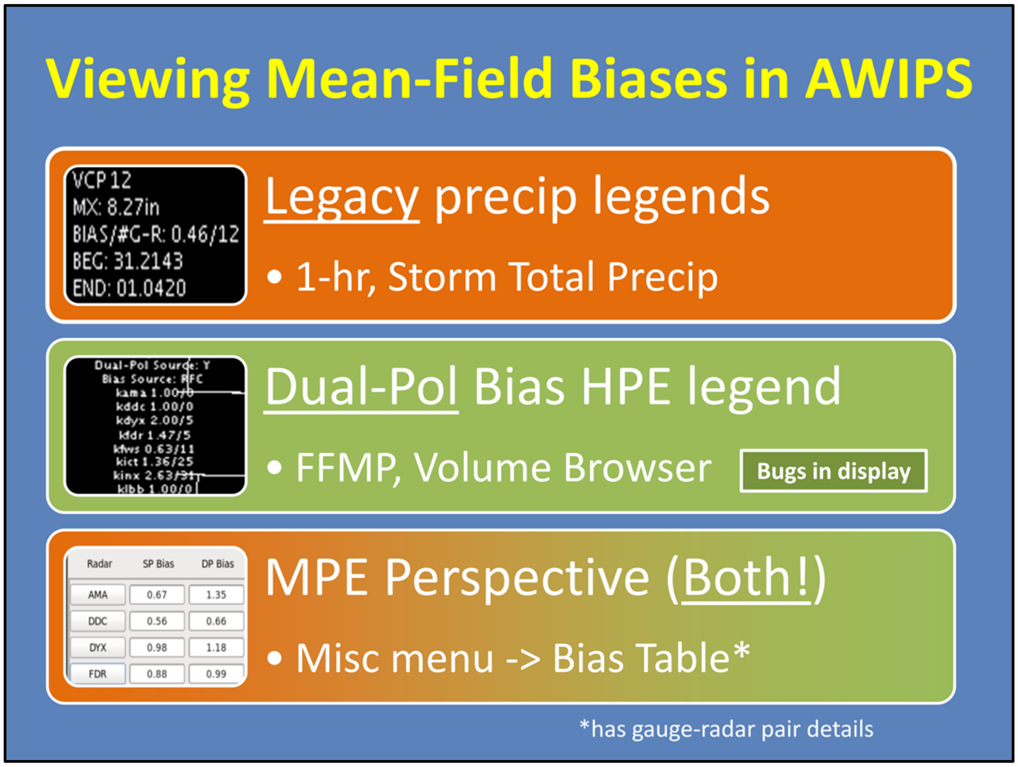

As diagramed on the previous slide, there are three ways to view mean-field bias information in AWIPS.

--For Legacy mean-field biases, look at the Legacy precip legends, as has been our focus throughout the lesson.
--For Dual-Pol mean-field biases, there's the Bias HPE product legend. This is a mosaic QPE product, and thus has Dual-Pol bias info for each radar contributing to the mosaic. It is an FFMP precip source, but is also accessed through the Volume Browser. However, there are some bugs in the display, as mentioned earlier.
--Finally, MPE has its own perspective in AWIPS, and there's a table forecasters can access with the bias info for both Legacy and Dual-Pol for all your radars, as well as details for the gauge-radar pairs. We don't expect you to go here during warning ops, but if you're interested in finding out how old your gauge-radar pairs are, this is where you'll find it. For more information on how to access this table, please see the supplemental slide at the end of this lesson.

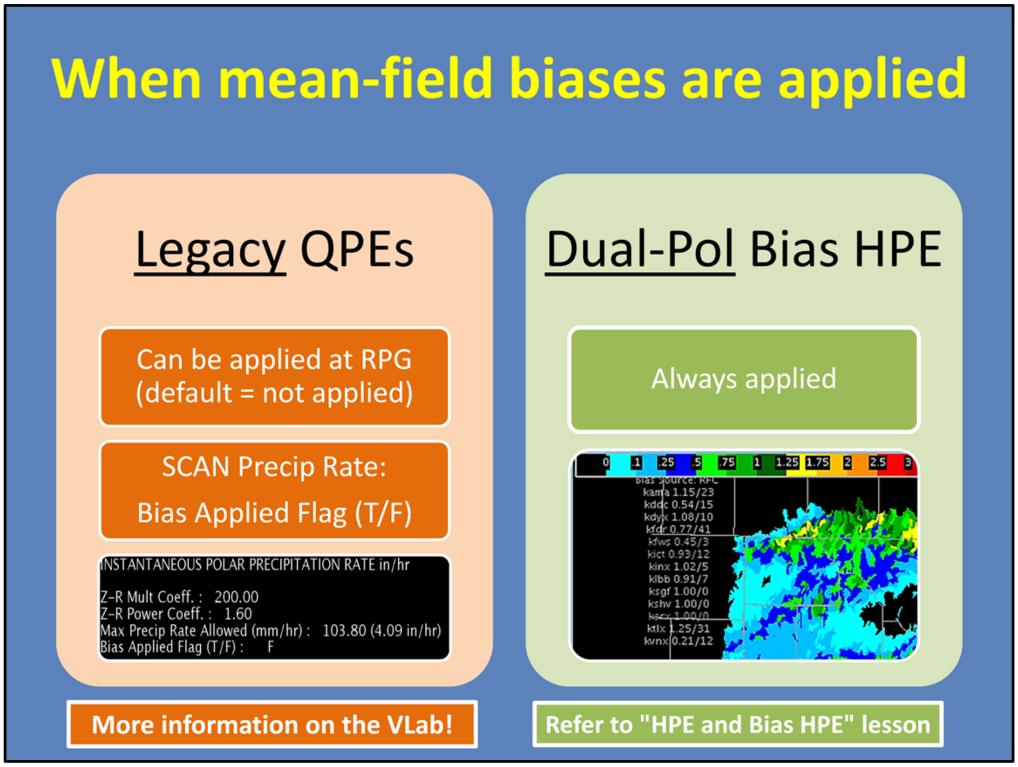

Even though there are several ways to *view* mean-field biases in AWIPS, there are only a couple places to apply them.

For Legacy precip estimates, the Legacy bias can be applied at the RPG, but the default is that it is not applied. To simply check whether it currently is or isn't, open the SCAN Precip Rate product for the radar of interest, and look at the legend shown here. The "Bias Applied Flag" will be false if it is not applied (the default), or true if it is. For more information on how to configure this flag to be viewable in the SCAN product and how to identify if the bias has been applied at the RPG, see the reference material on the VLab.

The other option is to use the Bias HPE product. This mosaicked Dual-Pol QPE product automatically applies each radar's Dual-Pol mean-field bias factor to the pixels within its coverage. There are more details on this product in the lesson "HPE and Bias HPE".

Dual-Pol Spatially-Varying Bias

QPE Source	Z-Rs	Bias Option
Legacy	**Single Z-R** (set at RPG)	**Mean-field**
		Spatially-varying (not commonly used)
Dual-Pol	**Spatially-varying** (based on precip type)	**Mean-field**
		Spatially-varying (available in Bias HPE; weighted based on distance from gauges)
MRMS	**Spatially-varying** (based on precip type)	**Spatially-varying** (weighted based on distance from gauges)

Until now, we've covered details related to mean-field bias options for Legacy and Dual-Pol. Before we jump into the MRMS biases, I wanted to briefly mention that Bias HPE also has the capability to apply a Dual-Pol spatially-varying bias instead of a mean-field bias. This requires a token change by your hydro focal point. I will save the details for the "HPE and Bias HPE" lesson, as well as the VLab reference material.

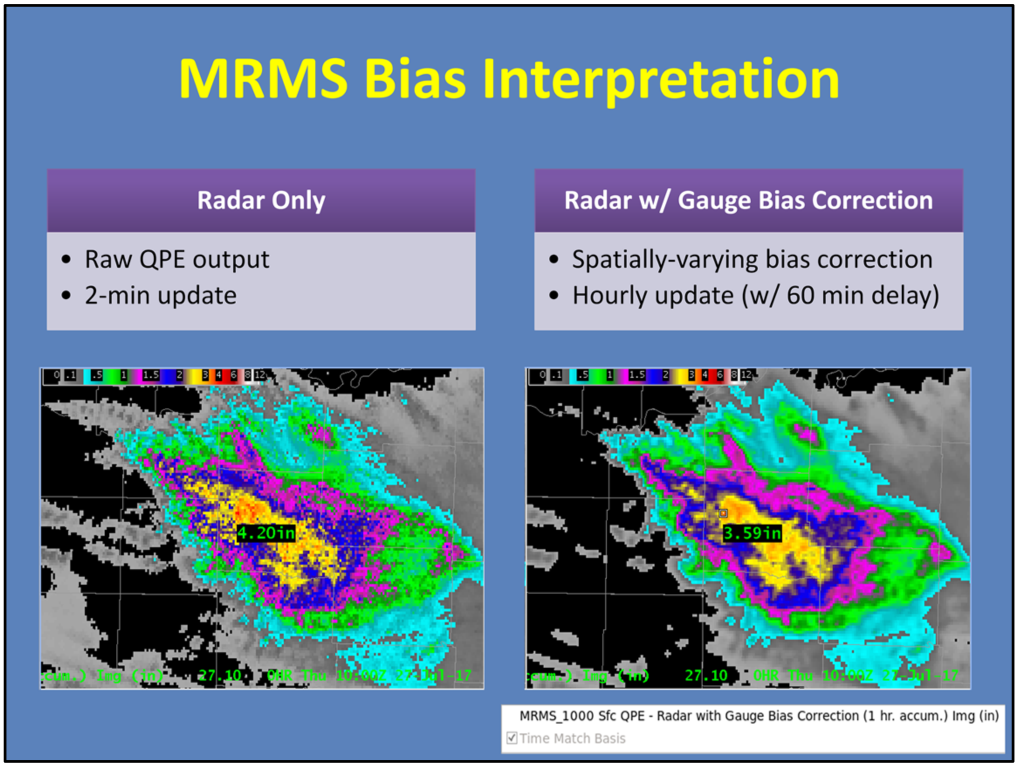

Alright, let’s switch gears to the MRMS bias interpretation, which is luckily much more straight forward.

To do it, you need to compare two products side-by-side. The first is the "Radar Only" product which tells you MRMS's raw QPE output. The second is the "Radar w/ Gauge Bias Correction" product. As the name implies, this product starts with the Radar Only output and then applies a spatially-varying bias correction based on gauges. Since it relies on gauges, this product is only created once every hour, and currently has a delay of about 60 minutes. To learn more, please visit the MRMS Products Course.

Even with this delay, there is still added value by doing this comparison. Load the products side-by-side in a 4-panel layout. Make sure the "Time Match Basis" is checkmarked on the *bias-corrected* product. This will force the Radar Only product to match the gauge product. Once they are time-matched, you get a clear visual comparison of the QPEs.

In this example, we see from sampling that the bias-corrected product tends to pull down the extremes where there are gauges. Where there are no gauges, the values remain unchanged.

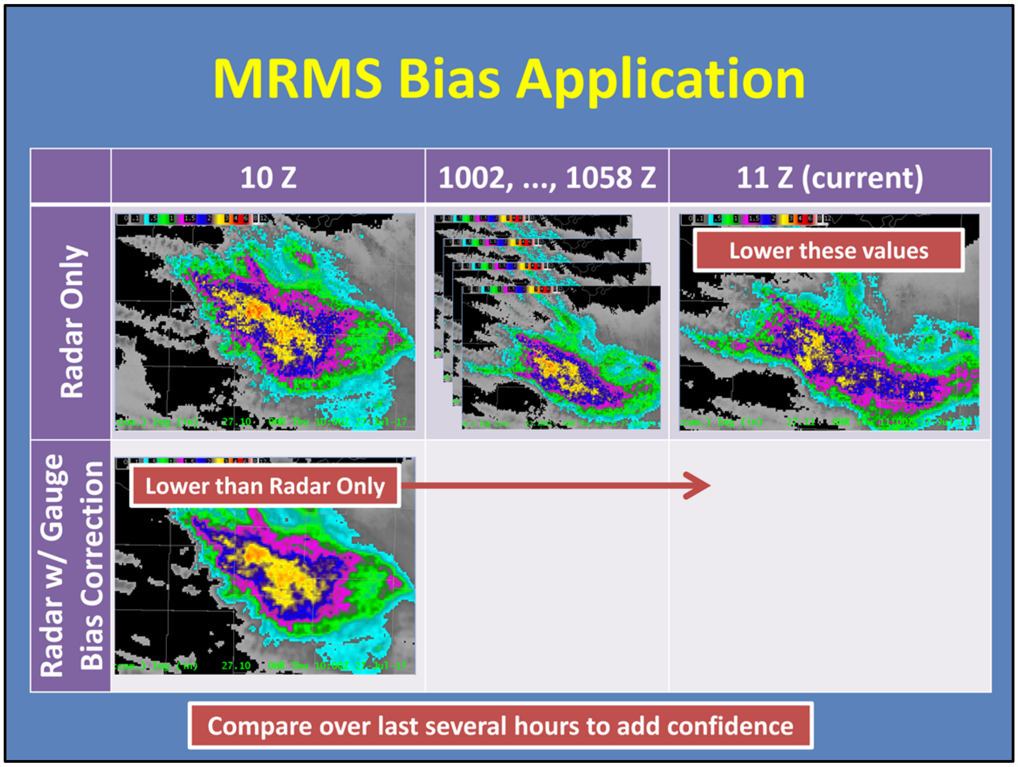

So how would I apply this in operations? Well, let's say current time is 11Z. Because of the gauge delay, the most recent bias-corrected product available is from 10Z. So I do the time-matched comparison at 10Z, sample out some values, and find that the gauges have been pulling down some of the raw radar estimates.

An hour's worth of Radar Only data has come in since this comparison time. So looking at my current Radar Only data at 11Z, I would mentally want to lower these values a bit, knowing what I analyzed from the previous hour.

Additionally, I would do this comparison over the last several hours to get a quick sense of whether there's a pattern. If there is, I have more confidence when extrapolating the current 11Z Radar Only values.

Back to Top

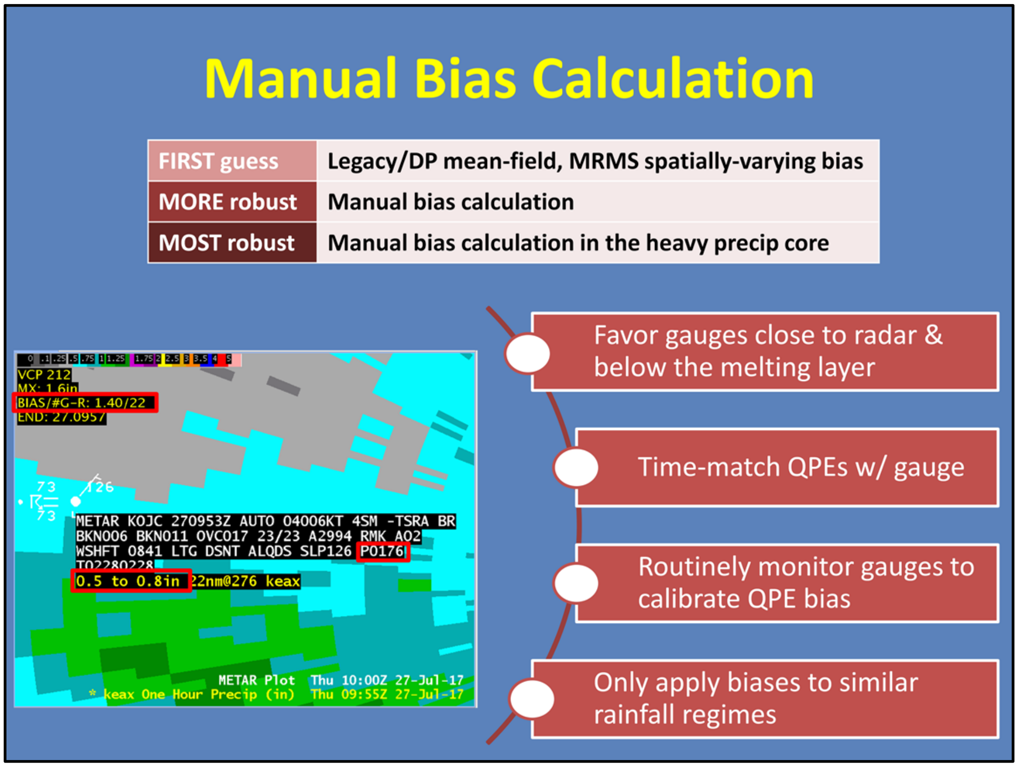

Mean-field biases and spatially-varying biases can be a good first guess for how the precip sources are doing for that event. However, a more robust method is to do a manual bias calculation using trusted gauges. This is especially true when the core of the heavy rainfall hits the gauge since it gives you ground truth in the location of the highest accumulations.

Take, for instance, this example that compares a METAR to the one-hour Legacy precip product. To start, notice that the Legacy mean-field bias is 1.40 based on 22 gauges. This is a first guess that Legacy is under-estimating. But what does the manual comparison say? The METAR observed 1.76 inches in the last hour (based on the P-group readout), while Legacy estimated 0.5 to 0.8 inches. This is a very clear under-estimate, with Legacy estimating less than half of what the gauge observed. This bias is fairly different than the mean-field bias, which tells you to multiply the QPEs by 1.4 to get them where they should be. In reality, the METAR comparison tells you to (at least) double them.

Here are some tips when doing a manual comparison with a gauge:
--For one, put more weight in gauges that are closer to the radar since the QPEs will be sampling lower in the atmosphere. Similarly, favor gauges that are below the melting layer. In this example, the gauge is 22 nautical miles away, which is good.
--Next, don't forget to time-match your QPE display with your gauge observations. Here, the METAR sent it's one-hour total at 0953, so I set the Legacy QPE display as close as possible which was 0955.
--Routinely monitor your gauges to calibrate your QPE bias. This could be every hour for METARs or checking in more frequently with Mesonets.

--Finally, only apply your biases to similar rainfall regimes. Don't expect a bias that you calculated in a convective core to necessarily hold true in a stratiform region of the storm. In this example, chances are the stratiform regions of this event will not need to be doubled since this is more of a convective signal.

Back to Top

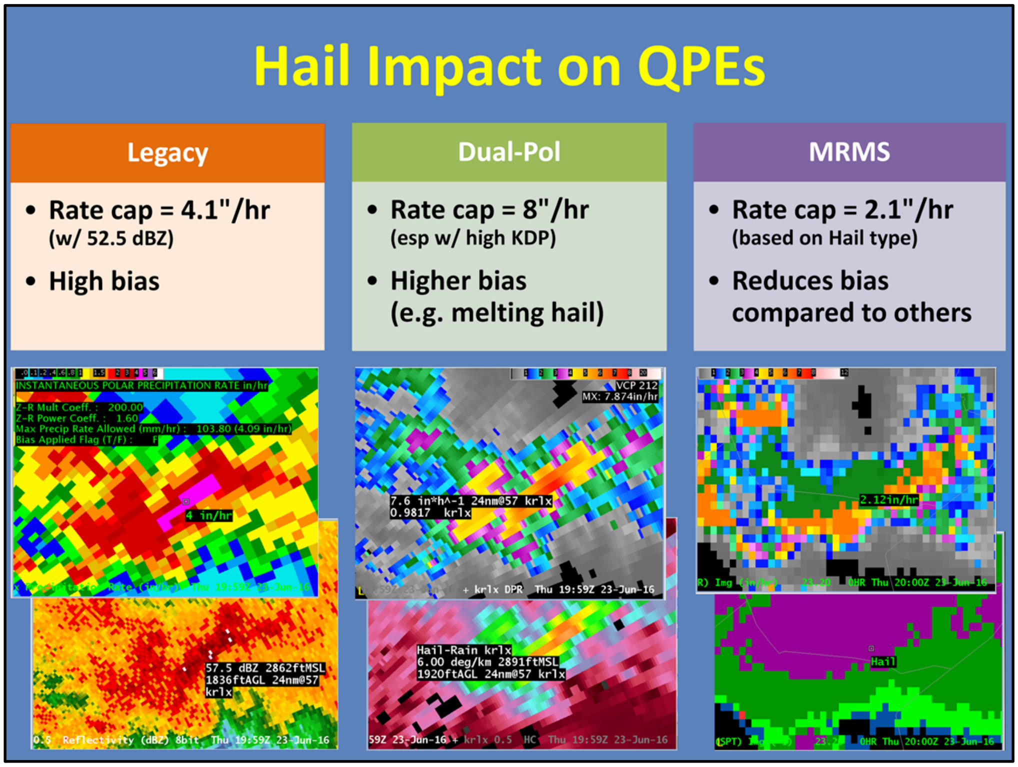

Lastly, I need to mention the effects that hail can have on doing bias calculations for QPEs. In general, hail can cause high biases for any precip source, so there have been adjustments made to each algorithm to try to account for it. The images below show the same storm for each precip source. Let's look at each one:

--For Legacy, there is a rain rate cap of 4.1 inches/hour. This is an adaptable parameter at the RPG, but recommended not to be changed. If you simply want to know your rate cap, look at the Instantaneous Precip Rate product legend, as shown here, which is loaded from the SCAN menu.

--For Dual-Pol, the rate cap is 8 inches/hour. This mostly occurs when melting hail causes KDP values to jump above about 4-5 degrees/km. To the radar, the melting hail "looks" like very large raindrops, so it allows for higher rates to occur, as seen in these images.

--Finally, MRMS bases its rain rate cap on precip type. For pixels deemed hail, the cap is 2.1 inches/hour. While this helps reduce the bias compared to the other two, the algorithm tends to over-assign the Hail class, as shown in purple. If you look at the reflectivity image on the far left, there is really only a minor area that should be tagged Hail. But MRMS broad brushes it, significantly capping the QPEs in those areas.

Back to Top

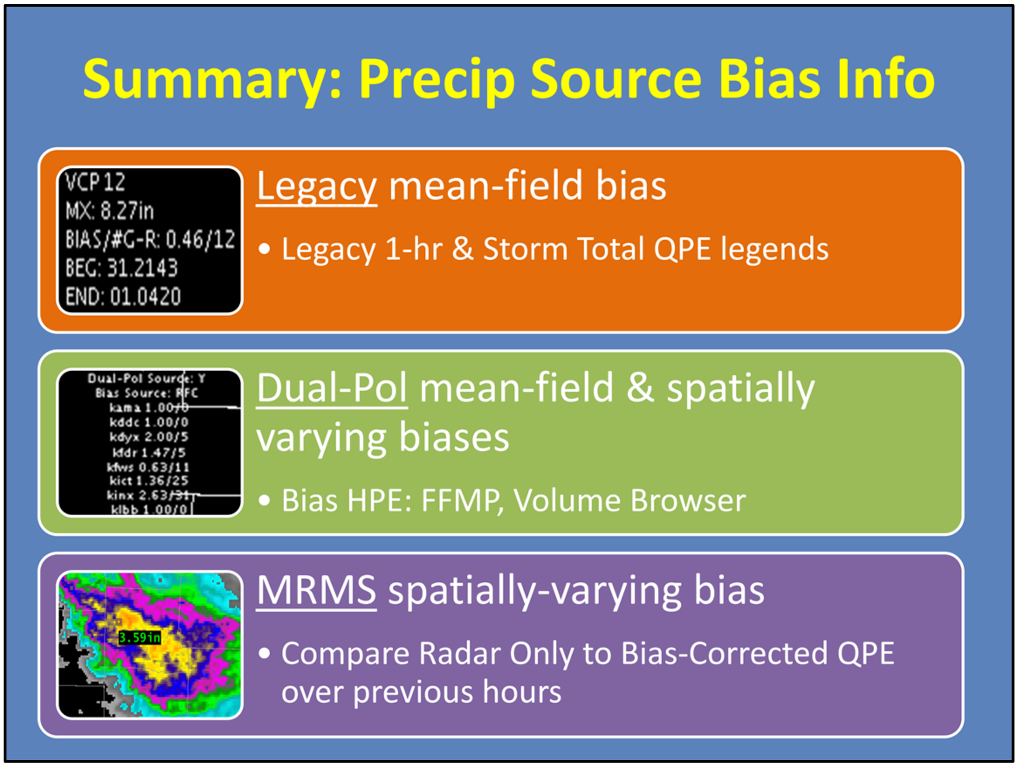

I know I threw a lot at you in this lesson, so here's a quick summary on where to look for each precip source's bias information.

For Legacy mean-field biases, simply look at the Legacy one-hour and storm-total QPE product legends.
For Dual-Pol mean-field biases and spatially-varying biases, use the Bias HPE product (either in FFMP or from the Volume Browser).
Finally, for MRMS spatially-varying biases, compare your Radar Only QPE to the bias-corrected QPE over the last several hours.

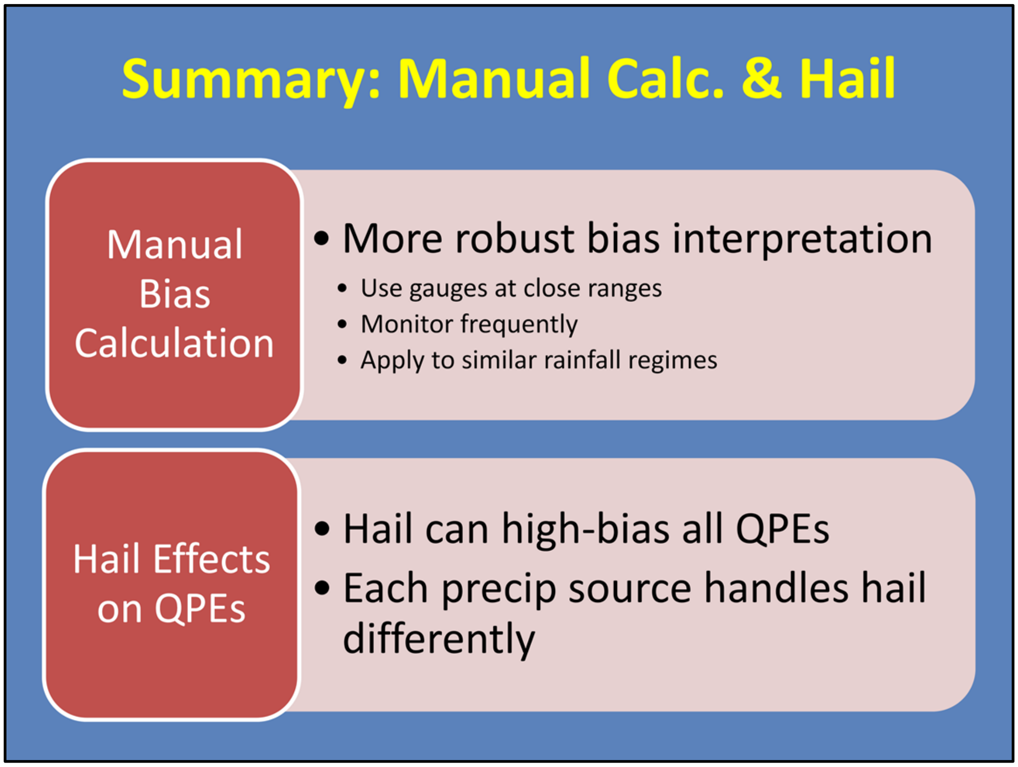

While mean-field and spatially-varying biases offer a good first guess, a manual bias calculation against a trusted gauge will usually provide a more robust bias interpretation. Performing this calculation using gauges at close range, monitoring at frequent intervals, and applying the result to similar rainfall regimes will further add confidence to your bias interpretation.

Additionally, always consider the effects that hail can have on QPEs, and look for these differences when comparing QPE sources.

This is the end of this lesson. When you are ready, please move onto the next slide to take the quiz and receive credit on the LMS.

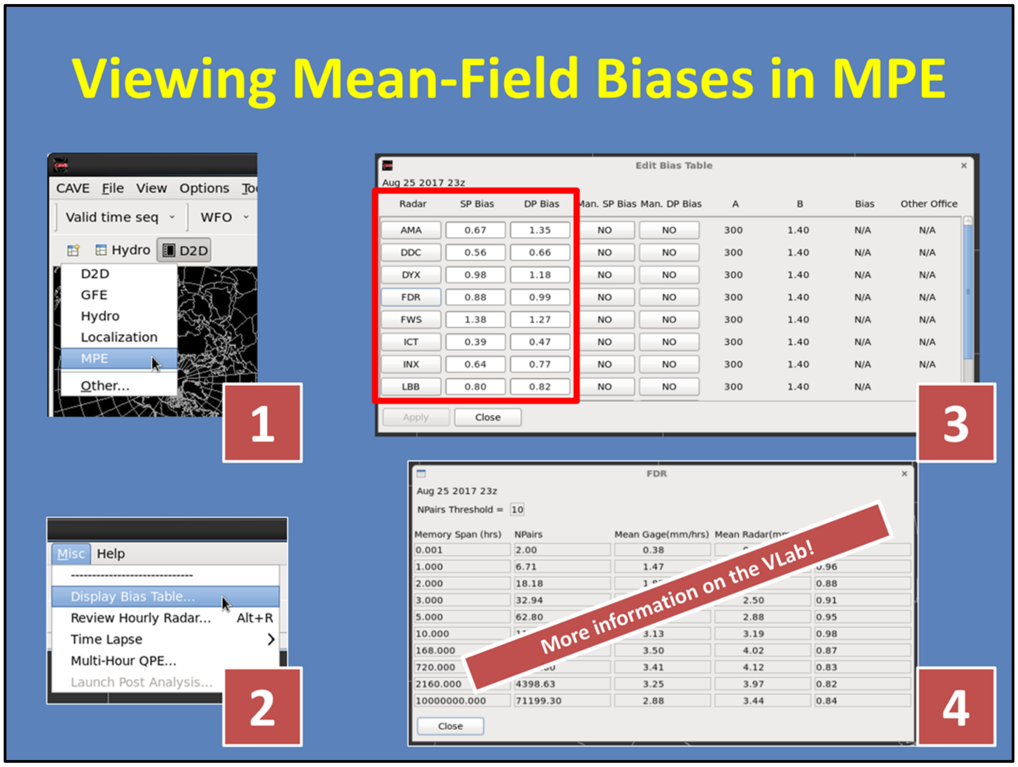

Here is your supplemental slide about viewing mean-field biases in MPE.

First, open the MPE perspective in CAVE. Once in MPE, navigate to the Miscellaneous menu, and click on the "Display Bias Table".

From this table, you will find all of the bias factors for both Legacy and Dual-Pol. Notice that each of the radar names is a clickable button.

If you clicked on one of these, you'll open a new sub-table with information about the gauge-radar pairs for that radar. Details on how to interpret that table are located at the VLab page in the Resources tab.

Back to Top

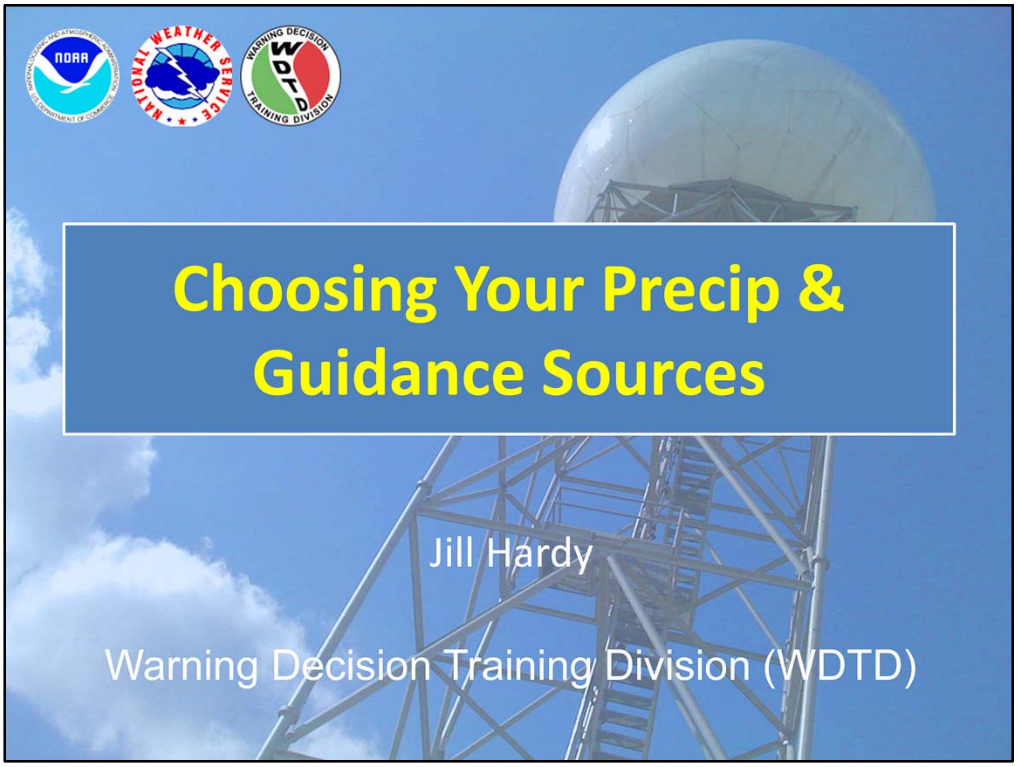

Hi, my name is Jill Hardy and welcome to this lesson on how to best choose your precip and guidance sources during the flash flood warning decision-making process.

Learning Objectives

- By the end of this lesson, you will be able to:
  - Identify basic approach to flash flood decision making
  - Identify characteristics of the five precipitation sources in AWIPS
  - Compare QPE to surface observations
  - Identify best practices when comparing QPEs to observations
  - Identify advantages of using FFMP for QPE-to-FFG comparison
  - Identify a FLASH product useful for assessing runoff
  - Identify caveats associated with QPE and FFG

Here are the learning objectives for this lesson. When you have finished reading them, please continue to the next slide.

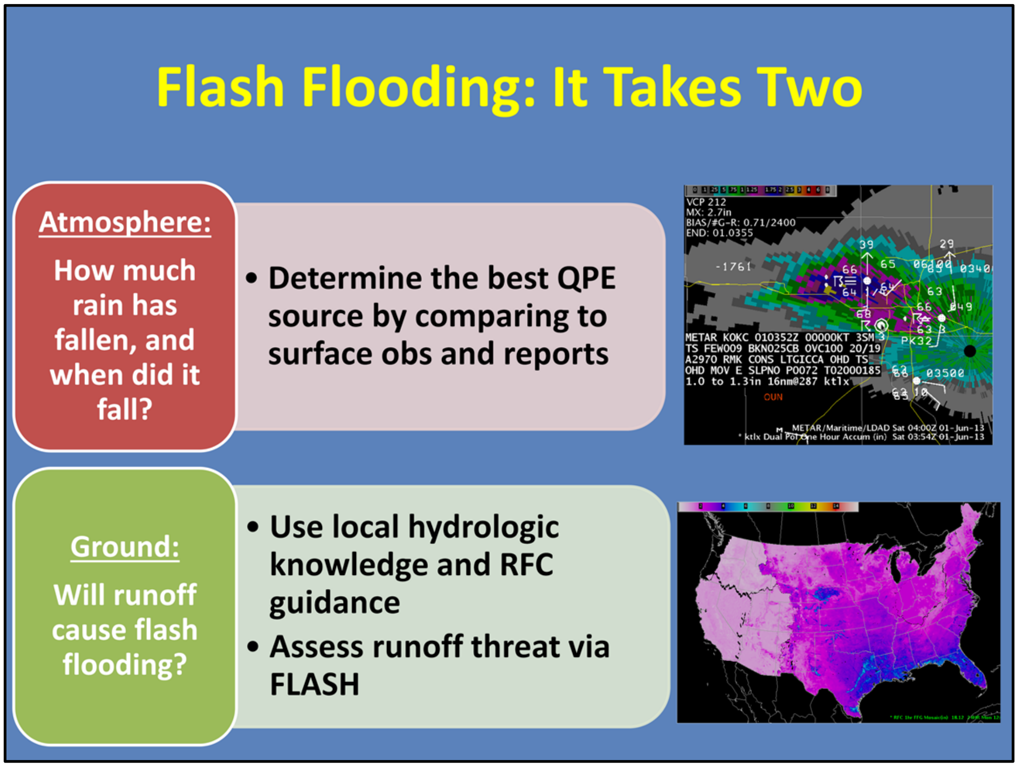

When considering a flash flood threat, there's really two questions to ask yourself. The first covers the atmosphere: "How much rain has fallen and when did it fall?". To do this, you will need to identify the QPE source that compares most favorably to surface observations and reports. There are multiple QPE options to choose from which will be covered in this lesson.

Next, consider the ground response. Ask yourself: "Will the runoff cause flash flooding?" Typically this starts with your local hydrologic knowledge and is supplemented with River Forecast Center (RFC) Flash Flood Guidance. There are a couple way to interpret this information, which we will discuss later. Additionally, you can assess the runoff threat using FLASH data.

Back to Top

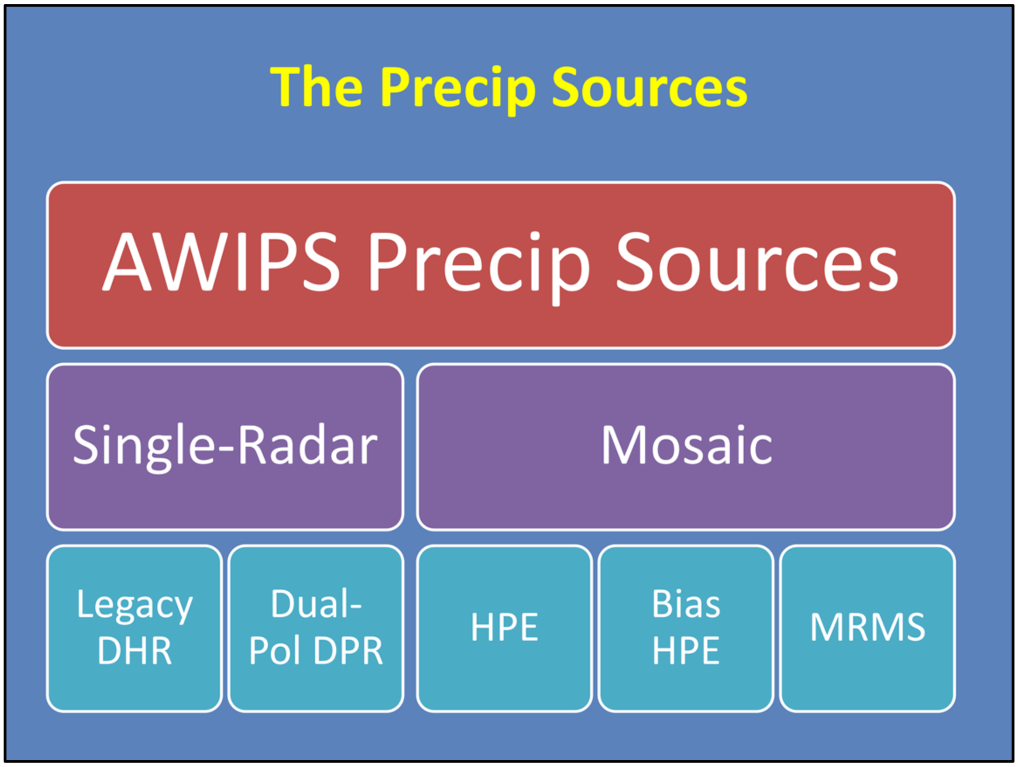

AWIPS has 5 precip sources available to choose from. For single-radar sources, there is the Legacy DHR and the Dual-Pol DPR. For mosaics, there's HPE and Bias HPE, as well as MRMS. Each one has strengths and limitations, so let's look at a detailed comparison of all of these on the next slide.

Choosing Initial Precip Sources

	Coverage	Dual-Pol?	Bias corrected?	Resolution	Accumulation products	Z-Rs
Legacy DHR	Single-radar	No	No (default, configurable at RPG)	1 km x 1 deg 3-6 min	1-, 3-, STP (and user-selectable)	Single Z-R, set at RPG
Dual-Pol DPR	Single-radar	Yes	No	0.25 km x 1 deg 3-6 min	1-, 3-, STA (and user-selectable)	Spatially-varying (based on HHC)
HPE mosaic	Mosaic	Yes	No	1 km x 1 km 5 min	1-hr	Spatially-varying (Inherited from DPR for each radar)
Bias HPE mosaic	Mosaic	Yes	Yes	1 km x 1 km 5 min	1-hr	Spatially-varying (Inherited from DPR for each radar, bias-corrected)
MRMS radar-only mosaic	Mosaic	No	No	1 km x 1 km 2 min	2-min update: 1-hr ≥ 1-hr update: all other accum	Spatially-varying (based on SPT)

There are many factors to consider when choosing a precip source to use during an event, specifically early in an event before you have many observations.

1. For one, consider the coverage area. The mosaics are good when you have multiple radars because you only have to open one grid or FFMP.
2. Next, is it Dual-Pol? Both HPE products default to using Dual-Pol precip sources. It is important to note that MRMS does NOT use Dual-Pol to calculate rain rates at this time, only for quality control purposes.
3. Do you want biases applied? Only Bias HPE has biases applied by default. However, DHR has the option available at the RPG. Note that MRMS does have bias-corrected products, but they are only updated at the top of the hour.
4. How do the resolutions compare? The single-radar DPR product has the best spatial resolution at 250 m by 1 deg. This is useful for isolated events that are close to the radar. However, MRMS has the best temporal resolution, at 2 minutes. This is useful for convective events with high rain rates.
5. Consider what accumulation products are available outside of FFMP. The single radar sources offer 3-hour and storm total accumulations. Additionally, they offer the ability to do one-time requests of user-selectable accumulations. On the other hand, the HPE products only produce 1-hour accumulations by default. And, MRMS only produces the 1-hour accumulation every two minutes, with all of the other accumulations being one-

hourly or longer.

6. Finally, how are rain rate relationships managed? DHR only has one Z-R, set at the RPG. Every other source uses spatially-varying Z-R relationships based on precip type.

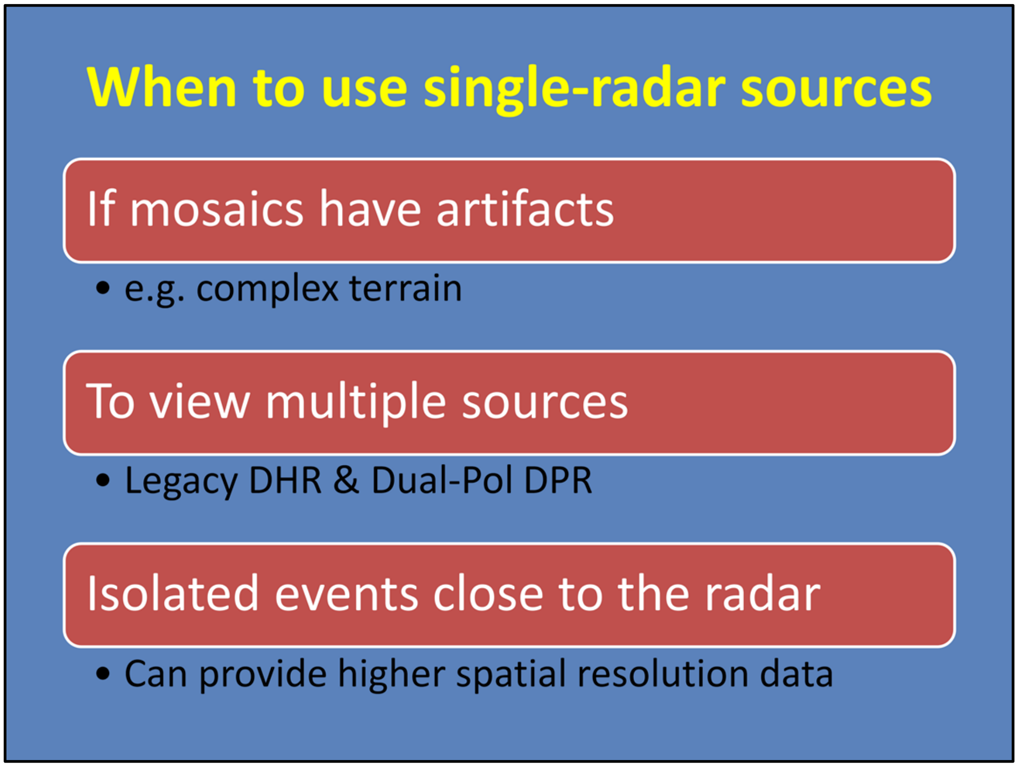

While generally it may seem that mosaicked products are preferred for their spatial coverage, there are times when we recommend using the single-radar products. They are good choices if the mosaics have artifacts, like when multiple radars are covering complex terrain. When you simply want to view more than one source, like looking at both Legacy and Dual-Pol. Or if there is an isolated event close to a radar, which may provide better spatial resolution data.

Choosing Initial Precip Sources

	Coverage	Dual-Pol?	Bias corrected?	Resolution	Accumulation products	Z-Rs
Legacy DHR	Single-radar	No	No (default, configurable at RPG)	1 km x 1 deg 3-6 min	1-, 3-, STP (and user-selectable)	Single Z-R, set at RPG
Dual-Pol DPR	Single-radar	Yes	No	0.25 km x 1 deg 3-6 min	1-, 3-, STA (and user-selectable)	Spatially-varying (based on HHC)
HPE mosaic	Mosaic	Yes	No	1 km x 1 km 5 min	1-hr	Spatially-varying (Inherited from DPR for each radar)
Bias HPE mosaic Bugs in bias factor display	Mosaic	Yes	Yes	1 km x 1 km 5 min	1-hr	Spatially-varying (Inherited from DPR for each radar, bias-corrected)
MRMS radar-only mosaic	Mosaic	No	No	1 km x 1 km 2 min	2-min update: 1-hr ≥ 1-hr update: all other accum	Spatially-varying (based on SPT)

So now, which precip source should you use? Really, the answer is all of them! It simply depends on which source seems to be capturing the nature of your current precip event. Recall from the lesson "Interpreting QPE Bias Information in AWIPS" that the choice comes down to the three overall precip sources: Legacy, Dual-Pol, and MRMS.

Generally speaking, Dual-Pol and MRMS are going to be your starting points because their rain rate relationships vary pixel-to-pixel based on precip class, which is meant to better capture the true nature of the precip occurring throughout the domain. The two sources do these calculations differently, so it's good to always look at both.

Dual-Pol then has 3 options, depending on your needs: the spatial resolution of the single-radar source, the coverage of the HPE mosaic, or the bias correction of the Bias HPE mosaic. Keep in mind that the Bias HPE product is buggy with how it displays the biases factors, so *only* use Bias HPE once you've taken our "HPE and Bias HPE" lesson and know how to work around this problem.

Of course, after your initial selection, routinely check these sources against observations to see if you need to make any changes.

Back to Top

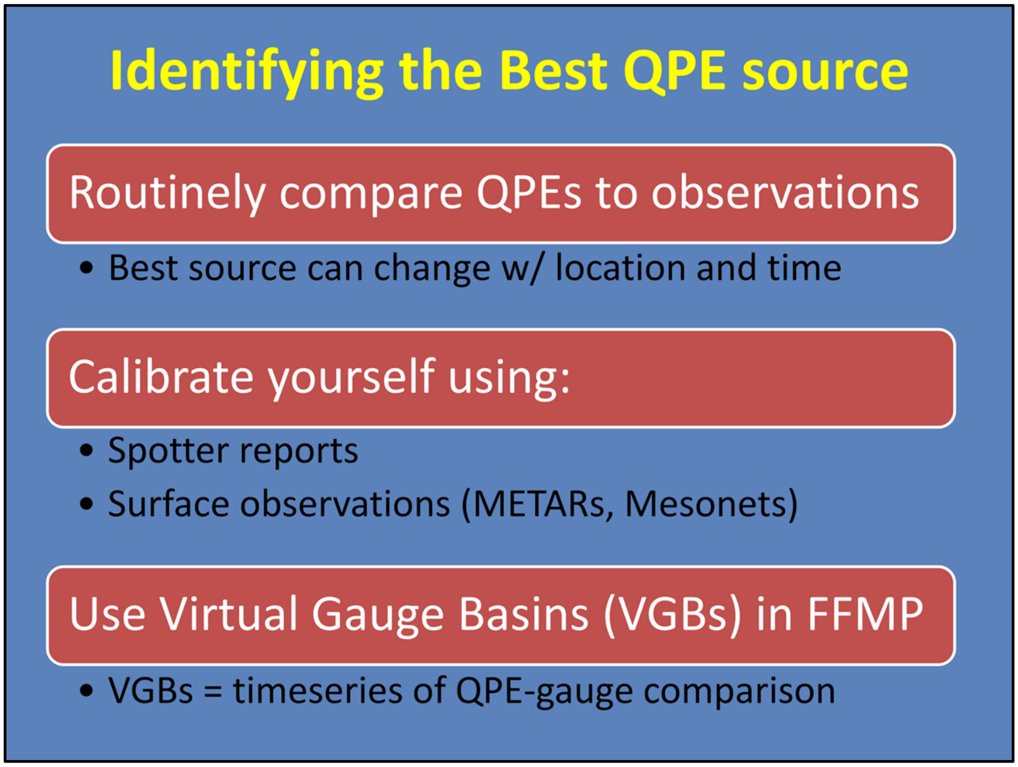

As the previous slide said, the best way to truly know which source is doing the best at any given time is to routinely compare them to observations throughout the event. It is not unusual for the optimal precip source to change across your CWA and with time.

Calibrate yourself using spotter reports, as well as sampling surface observations, such as METARs, Mesonet stations, and CoCoRaHS,.

Additionally, the Virtual Gauge Basins (or VGBs) feature in FFMP is also a convenient way to compare a QPE source to observations. This process is a little more involved, so refer to the lesson in WOC Flash Flood to learn more.

Back to Top

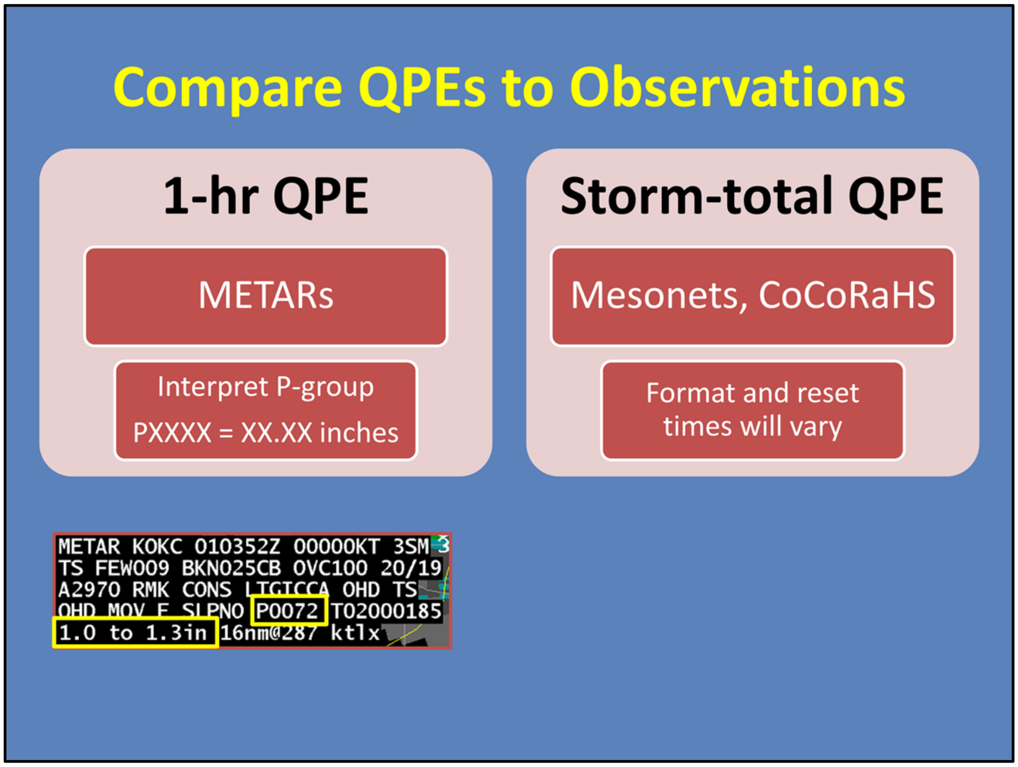

Manually comparing QPEs to trusted gauges is a robust method for determining how your precip sources are doing for that event.

For one, compare one-hour totals to METARs, since METARs report hourly. To find the precip total, look for the P-group. The 4 digits correspond to rainfall in inches, using this conversion.

Here is an example. We see the METAR estimates 0.72 inches, while the 1-hour radar QPE estimates 1.0 to 1.3 inches. This shows that the precip source is overestimating by a quarter to half an inch. Now you can use this information to self-calibrate the QPE estimates.

For storm total QPEs, compare to local Mesonet and CoCoRaHS stations if you've got them. Keep in mind that these networks format their outputs different, as well as reset at different times. For instance, one network may reset at midnight local or 00Z. So you may need to do some on-the-fly calculations if this happens.

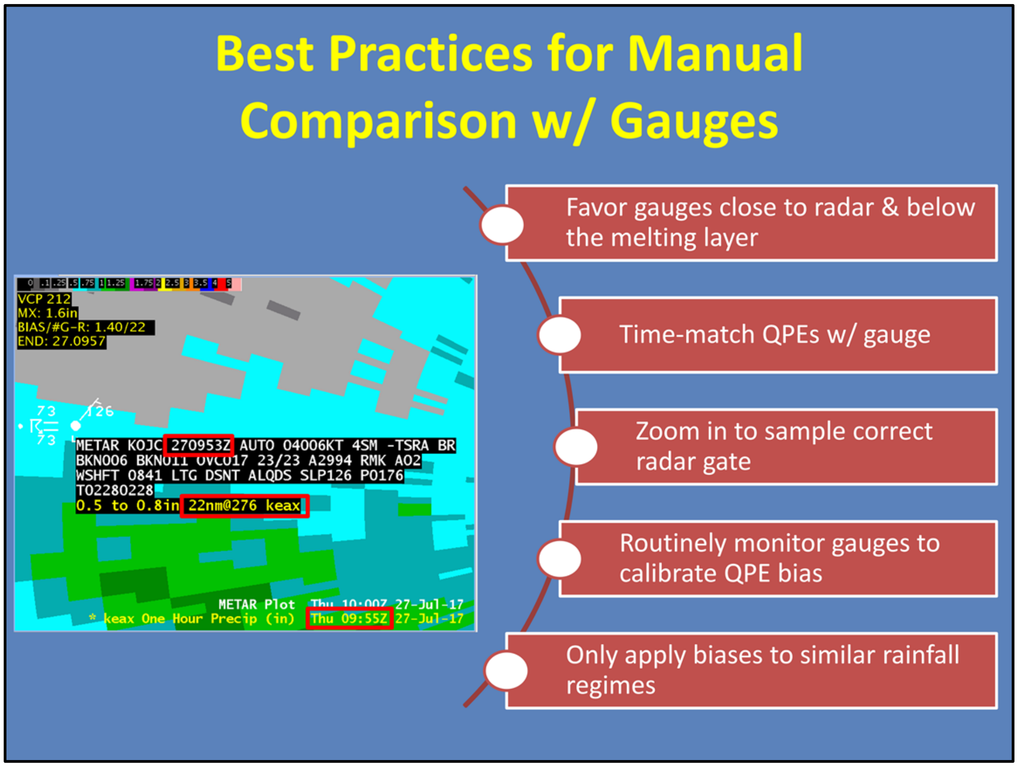

Here are some tips when doing a manual comparison with a gauge:
--For one, put more weight in gauges that are closer to the radar since the QPEs will be sampling lower in the atmosphere. Similarly, favor gauges that are below the melting layer. In this example, the gauge is 22 nautical miles away, which is good.
--Next, don't forget to time-match your QPE display with your gauge observations. Here, the METAR sent it's one-hour total at 0953, so I set the Legacy QPE display as close as possible which was 0955.
--Also, don't forget to zoom all the way in when you do these comparisons, to ensure you've got the right radar gate.
--Routinely monitor your gauges to calibrate your QPE bias. This could be every hour for METARs or checking in more frequently with Mesonets.
--Finally, only apply your biases to similar rainfall regimes. Don't expect a bias that you calculated in a convective core to necessarily hold true in a stratiform region of the storm.

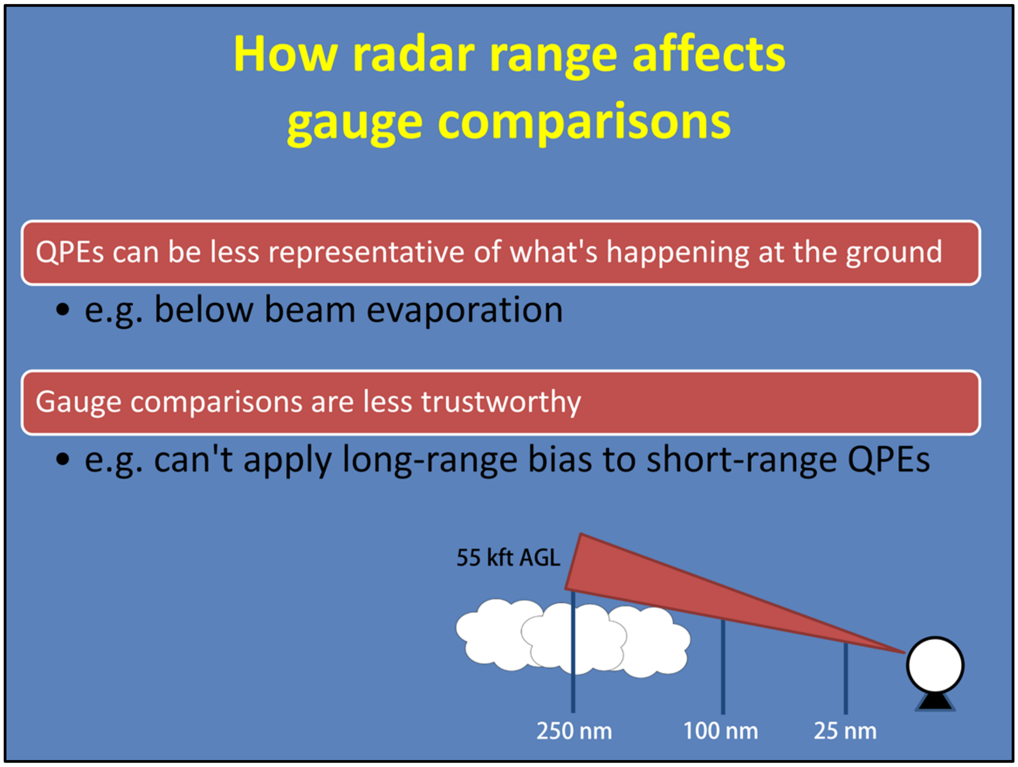

We recommend that you focus on doing these surface obs comparisons at closer ranges because, at long ranges, there are two things to think about:

1) Precip estimates can be a lot less representative of what's happening at low levels. Consider the depiction below, where a radar estimate at 250 miles away could be sampling as high as 55 kft above ground level.
2) Because of this, it makes comparisons with ground observations less trustworthy. For instance, if your gauge networks are all really far from the radar, be wary of your bias information. You can't apply a bias from long range to radar data at short range because they are sampling different parts of the atmosphere.

Back to Top

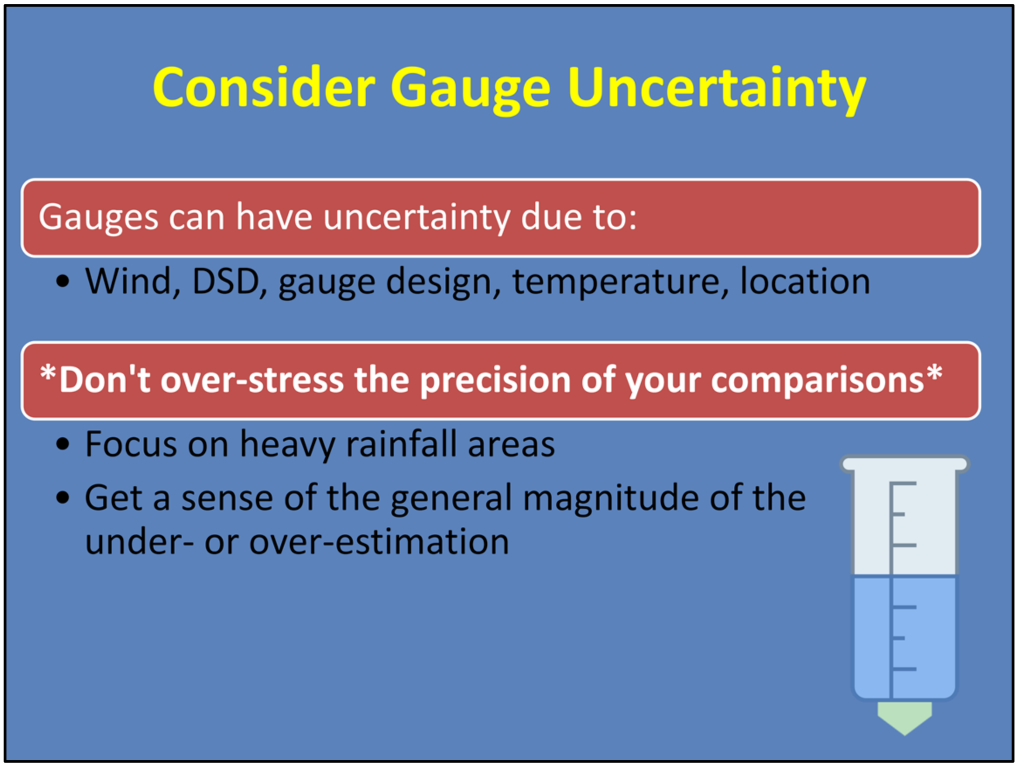

A quick note: Just as we know there's uncertainty with radar estimates of rainfall, there is also uncertainty with gauge estimates of rainfall. This is outside the scope of this lesson, but keep it in mind when doing your comparisons. Gauges can underestimate due to factors like wind, drop size distribution in storms, the mechanisms of the gauge itself, temperature, and location of the gauge.

From the COMET module "Rain Gauges: Are They Really Ground Truth?", it's within reason to see anywhere from 10-25% variation of gauge totals, simply due to the factors listed above.

So the takeaway here is this: Don't dive too deep into the precision of your QPE-to-gauge comparisons. Focus on areas receiving heavy rainfall to do your comparisons, since these areas will have the highest impact on flash flooding. Yes, some of the variation can be attributed to uncertainties, but getting the general magnitude of whether you're off by let's say 0.25-0.5 inches or 1.0-2.0 inches is useful in your decision-making.

Back to Top

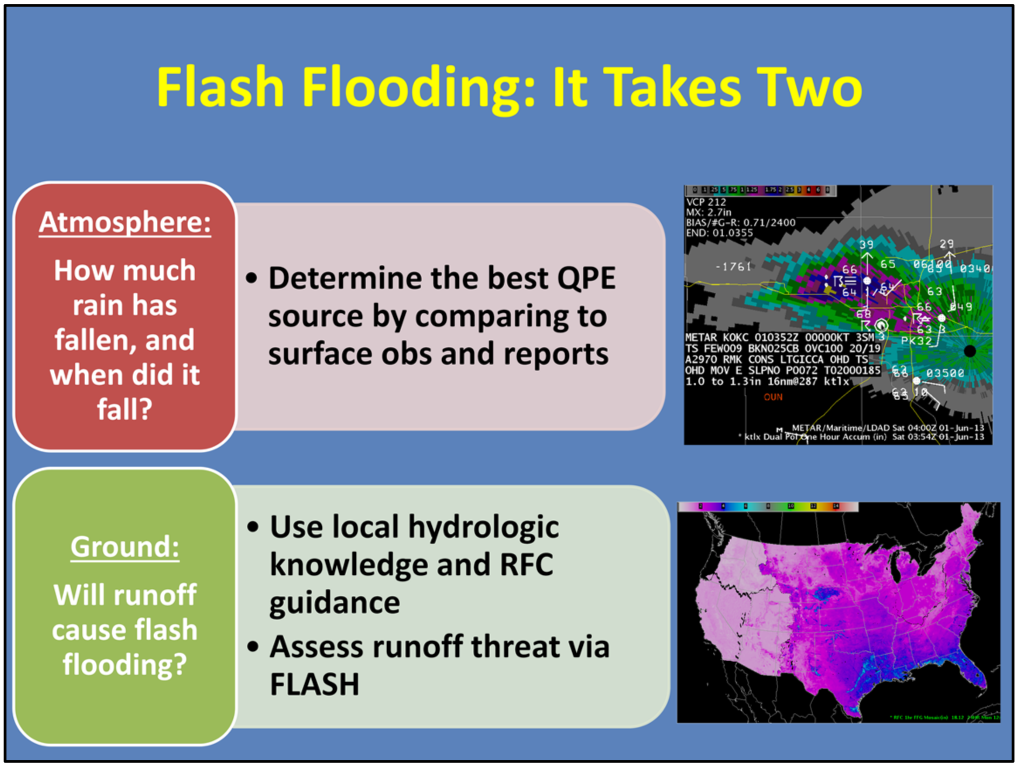

Okay, so you've identified the best QPE. The next step is to determine if the runoff will cause flash flooding. This is where you use your local rules of thumb for flashy basins, compare the best QPE to the RFC Flash Flood Guidance, and assess runoff threat via FLASH.

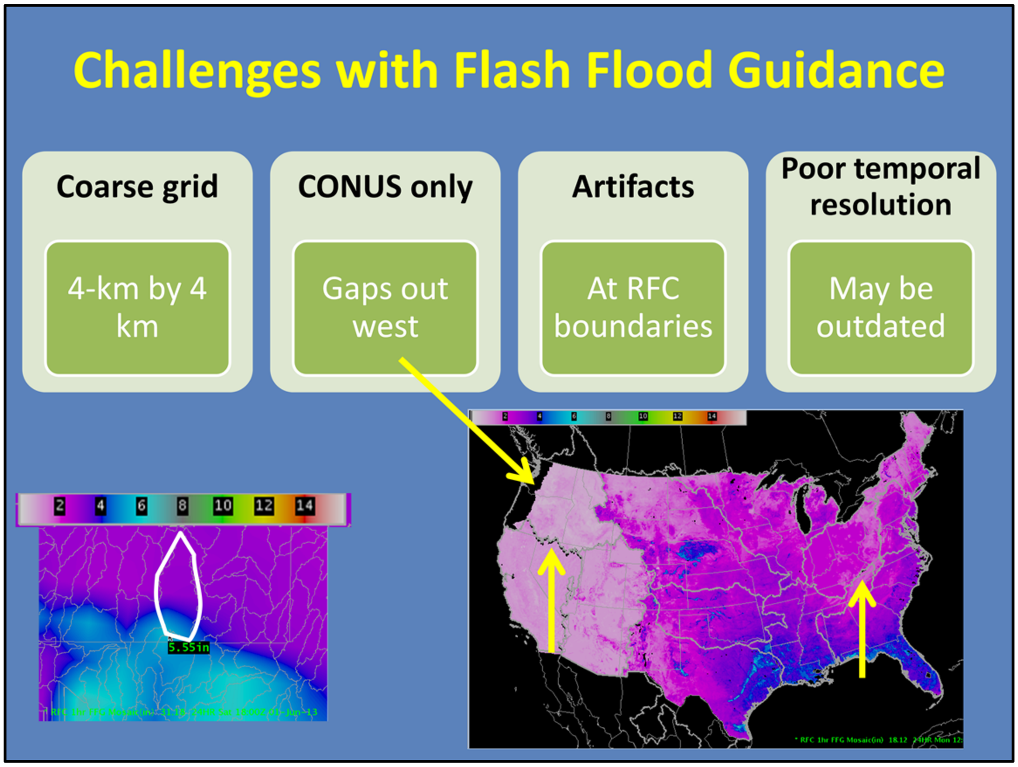

If you have an idea of how much rain has fallen and you are going to compare it against RFC flash flood guidance, it is important to note some of the challenges.

1. For one, RFC FFG has a rather coarse grid (approximately 4km by 4km) relative to the small size of many basins. Take this zoomed in 1-hr FFG product. Over this particular basin, FFG ranges from a little over 2 inches to the north to 5.55 inches to the south. FFMP is going to average this out to one number, which may not adequately represent the hydrology over this basin.
2. Second, it's coverage is only over the CONUS, with some gaps out west, as shown here.
3. Next, there are artifacts along some RFC boundaries where different methods of calculating guidance result in non-realistic sharp gradients, which we can see when overlaying the boundaries.
4. Finally, FFG is only updated up to 4 times a day. Oftentimes with flash flooding, you may have a fast-moving, high-rate event that will saturate the ground. And before FFG has a chance to update, another event moves over the same area. FFG's poor temporal resolution could inhibit your interpretation if you don't take into account the earlier storm.

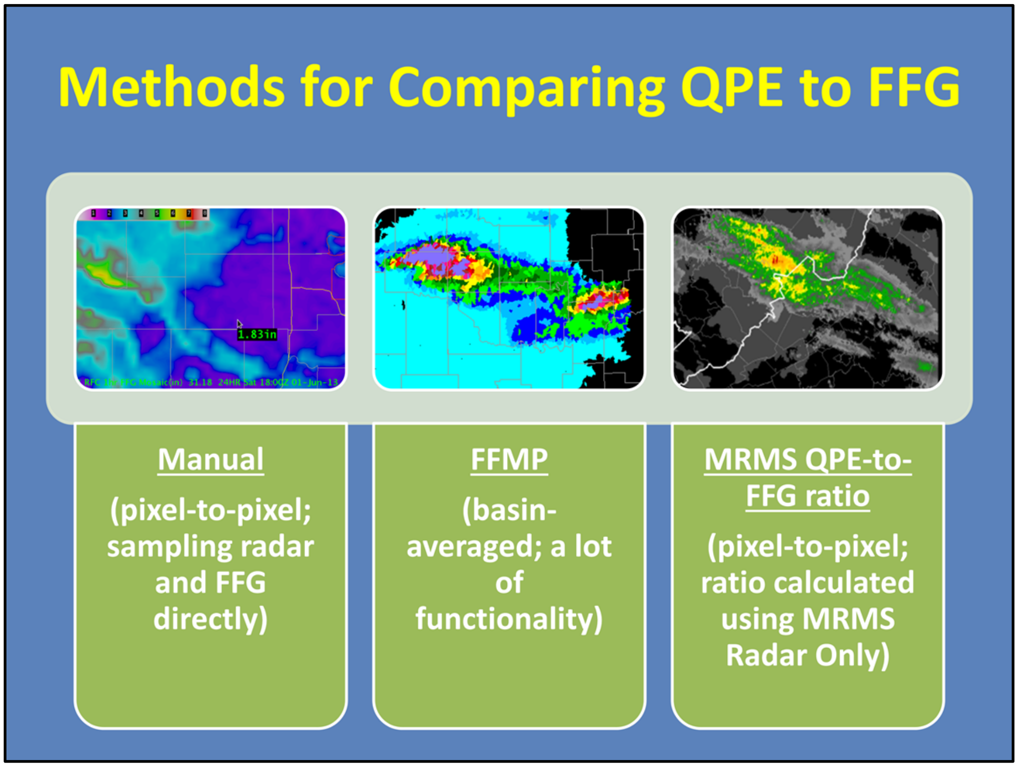

In AWIPS, there are three methods to compare QPE to Flash Flood Guidance. For one, you can manually load the radar accumulations and RFC FFG, and sample out the values directly. The next option is to use FFMP, which provides basin-averaged comparisons. FFMP has a lot of functionality that we will discuss shortly. Finally, there's the MRMS QPE-to-FFG ratio product. As it sounds, it provides a ratio comparison based on the MRMS precip source.

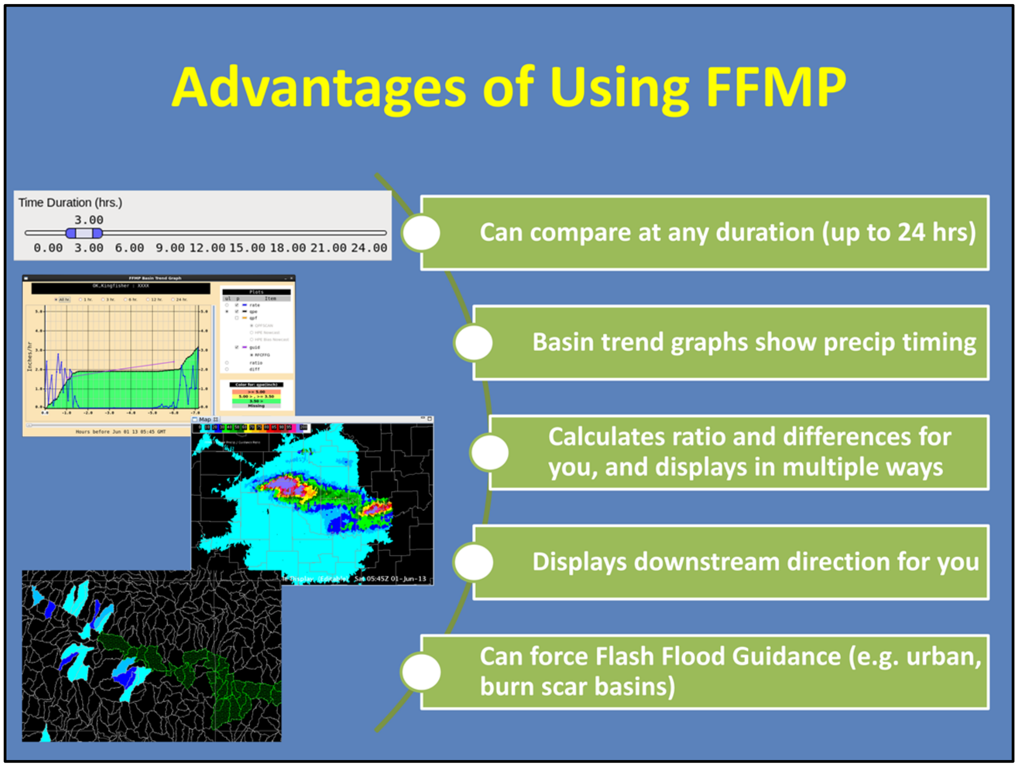

When it comes to comparing QPEs to Flash Flood Guidance, FFMP is your most powerful tool. Here are some advantages to using it:

1. First, FFMP accumulates precip every time a rate product is ingested, so one of its greatest strengths is the ability to display any accumulation duration (up to 24 hours).
2. This accumulation-on-the-fly approach also allows FFMP to display the precip timing information in the basin trend graphs. This is helpful to determine when precip occurred with training storms over long periods of time.
3. Another fundamental strength of FFMP is that it calculates QPE and FFG ratios and differences for you, and can display this information in a number of ways.
4. FFMP will also show you drainage and downstream information quickly and easily.
5. Finally, FFMP also has a tool to create your own forced FFG which can be very useful in urban and burn scar basins.

Back to Top

MRMS QPE-to-FFG Ratio

Analogous to ratios in FFMP, just not basin-averaged

- Pro: Do not smooth out extremes
- Con: Do not account for hydrologic extent (i.e. basins)

Calculated using nationally-mosaicked FFG grid

- Does not account for local variations (i.e. forced FFG)

Use when:

- You want a quick look at ratios
- You want to easily overlay with other products

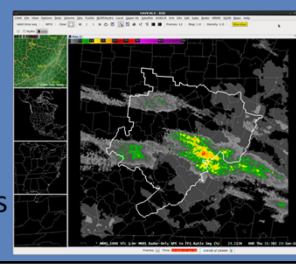

There are more details in the "FLASH Best Practices" lesson in the WOC Flash Flood course, but here is a quick introduction to this ratio product:

It is analogous to interpreting ratio in FFMP, it's just not basin-averaged. A benefit to looking at ratios in this way is that it does not smooth out the extremes which may help you pinpoint higher threat areas. However, it also doesn't account for the hydrologic extent of the threat the way basin-averaging is made to do.

Another caveat is that the RFC Flash Flood Guidance grid that is used here is from a national mosaic created hourly at WPC. Therefore, any local changes made to the grid, like forced FFG, are not captured.

To summarize, this product is useful when you want a quick look at ratios without opening FFMP, or when you want to easily overlay the data with other products. Otherwise, FFMP is far more customizable.

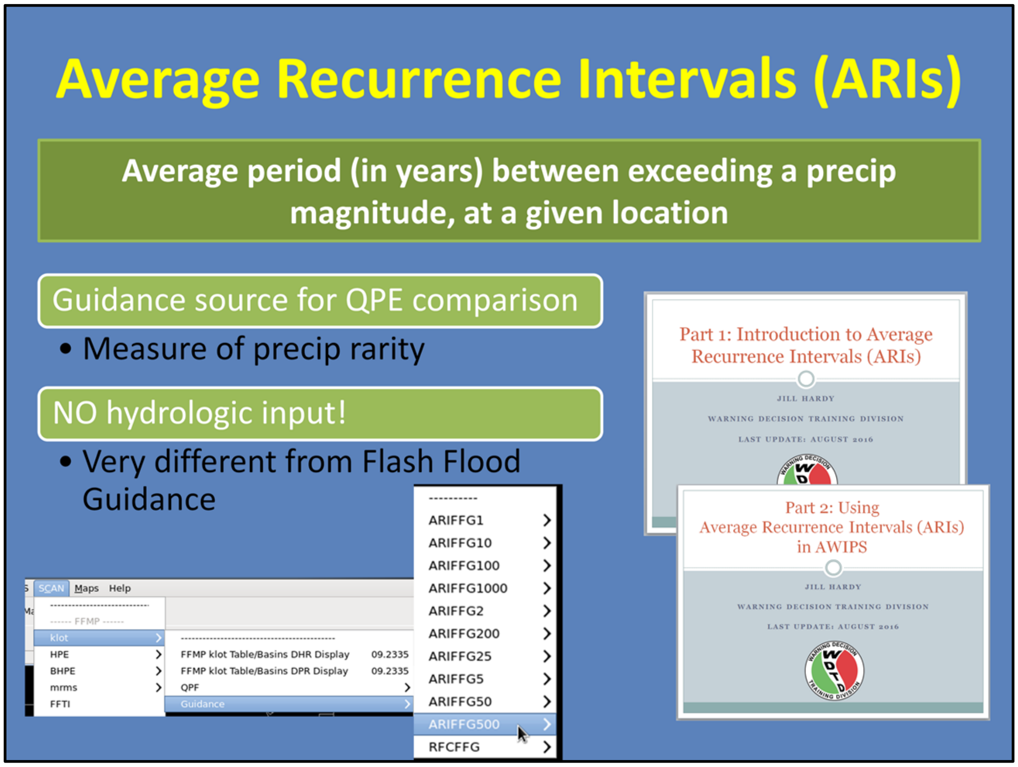

Since 2016, AWIPS includes Average Recurrence Intervals (or ARIs). An ARI is defined as the average period (in years) between exceeding a precip magnitude, at a given location. You are probably more familiar hearing ARIs used like “Yesterday’s 24-hour rainfall total was a 100-year rainfall event”.

We bring up ARIs here because they are a guidance source for QPE comparison, and you will see them when maneuvering around AWIPS and FFMP. For instance, here is a SCAN FFMP menu. Under the Guidance submenu, it’s no longer just RFC FFG, but also a lot of ARI data, as well. One of the most important things to remember when using this dataset is that it was created solely as a measure of precip rarity. It does not include any hydrologic inputs. Therefore, it is fundamentally different than Flash Flood Guidance, even though they appear in a similar fashion throughout AWIPS and FFMP.

Before you begin using this dataset, we highly recommend taking our two ARI lessons. There are a lot of details covered in these lessons that are crucial to effectively interpret and relay this information. These lessons are included in WOC Flash Flood, if you plan to enroll in this course.

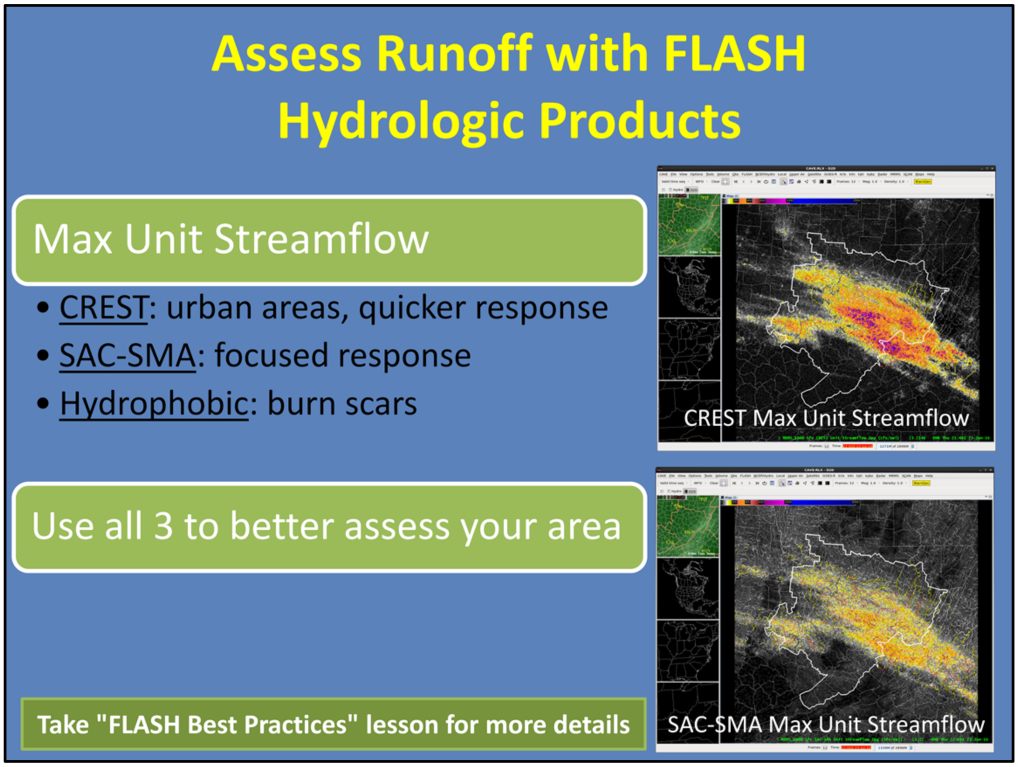

Last, but certainly not least, is to assess the runoff threat using the FLASH hydrologic products.

Probably the most useful product is the Max Unit Streamflow since it highlights areas experiencing higher than normal flow. There are 3 models to consider. CREST is the top image here, and tends to better highlight urban areas, as well as having an overall stronger, quicker response compared to SAC-SMA (which is the bottom image). SAC-SMA has a more focused response, and better highlights specific basins. Finally, the hydrophobic is a good proxy for burn scars.

The majority of research and evaluation has been on the CREST model, so we have the best understanding of its applications. However, use all 3 in operations to get a better idea of what works for your area. And don't forget that FLASH has other products that can also help in warning operations, so take the "FLASH Best Practices" lesson for more details.

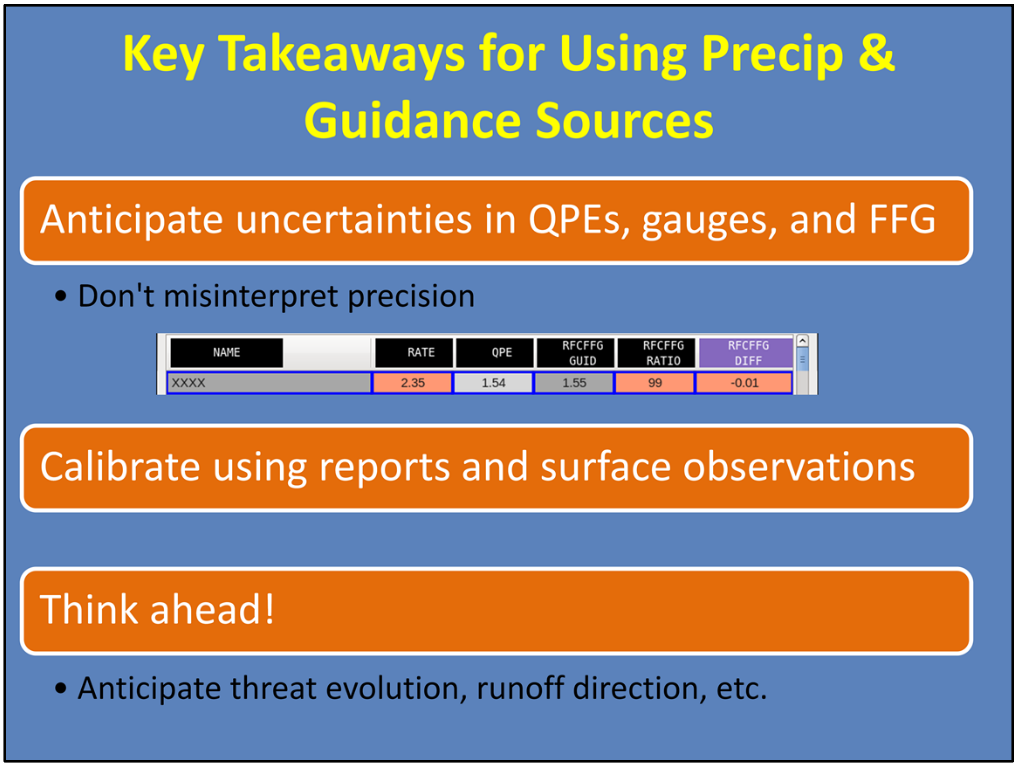

There are several key takeaways for using your precip and guidance sources.

For one, anticipate uncertainties in QPEs, gauges, and FFG. It is not uncommon to encounter uncertainties on the order of 25%, or more. However, AWIPS generally displays these values with two decimal places. Don't misinterpret this precision. For instance while the selected basin may appear 0.01" below FFG, this could easily be a quarter inch above or below FFG due to uncertainties in both the raw QPE or FFG data.

Therefore it is important for you to be routinely calibrating QPE using reports and surface observations, keeping in mind that surface obs can also have their own uncertainty.

Finally, always think ahead. It is easy to become fixated on the complexities of what is going on now with tools like FFMP. Anticipate threat evolution by considering where the storms are moving and what the hydrological conditions will be in those areas. This will give you important lead time when drawing your FFW polygons.

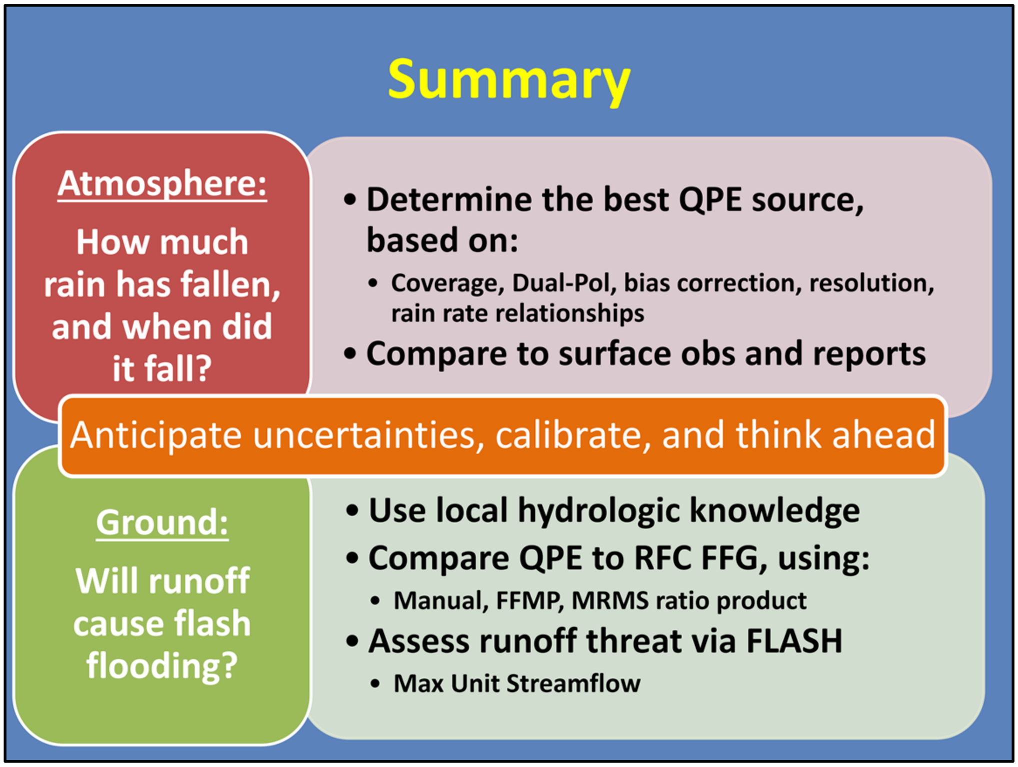

To summarize, the basic approach to flash flood decision making begins by assessing how much rain has fallen, and when. To do this you need to evaluate multiple precip sources and choose the best precip source based off factors like: the coverage of the product, whether it's Dual-Pol or has a bias correction, the resolution, and how the rain rate relationships are calculated.

Just as important as your initial precip source selection, is to routinely compare all precip sources with surface observations, gauges, and spotter reports. The best precip source can change over the course of an event.

The next step determines if the runoff will result in flash flooding. Use your local hydrological knowledge. Then compare your QPE to the RFC FFG. Consider the challenges that come with using FFG, as well as the various methods to interpret it. Additionally, assess runoff threat using the FLASH hydrologic products, particularly the Max Unit Streamflow.

Finally, always anticipate uncertainty in QPEs, gauges, and FFG. Calibrate accordingly, and remember to think ahead to anticipate future risk.

This is the end of this lesson. When you are ready, move onto the next slide to take the quiz and receive credit on the LMS.

Hi, my name is Jill Hardy and welcome to this lesson which will focus on using FFMP to diagnose a flash flood threat.

Back to Top

Learning Objectives

- By the end of this lesson, you will be able to:
  - Identify when to use All and Only Small Stream Basins layer versus County layer
  - Identify why QPE, Ratio, and Diff are useful for flash flood decision-making
  - Interpret QPE, Ratio, and Diff in FFMP
  - Interpret the Basin Trend Graph, specifically the all-hours graph
  - Identify when to use downstream trace in FFMP in warning decision making

Here are the learning objectives for this lesson. When you have finished reading them, please continue to the next slide.

Back to Top

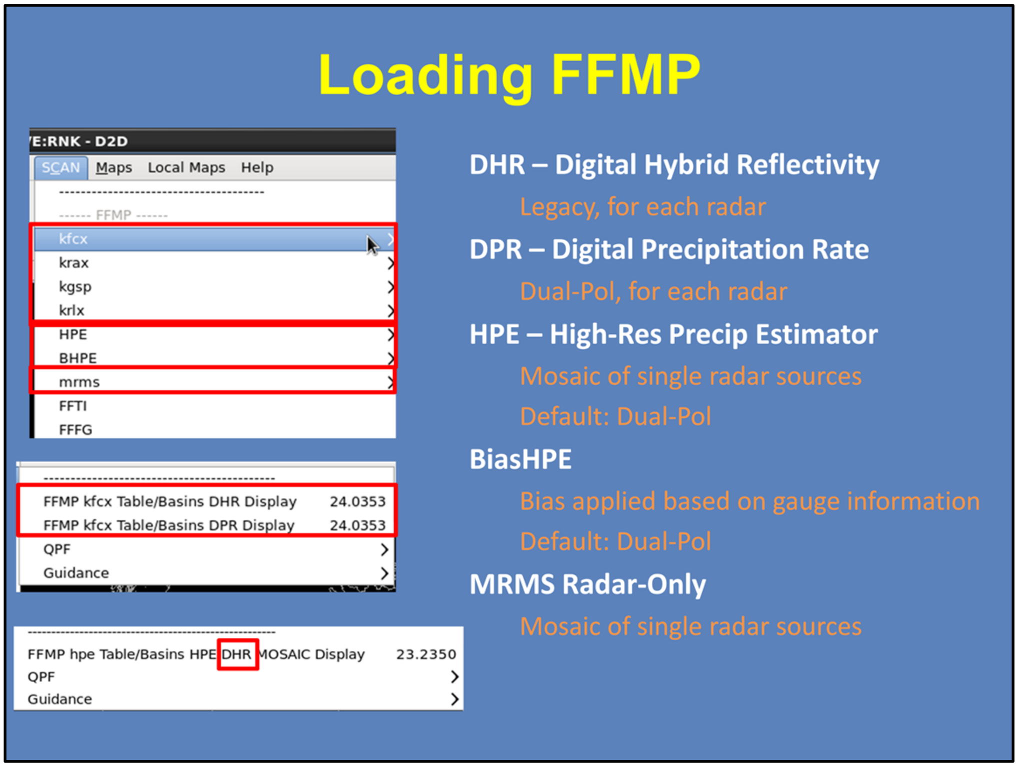

First thing's first: loading FFMP with your desired precip source from the SCAN menu.

The single-radar products are available under the menu referenced by the radar name. As seen here, under each radar submenu, there is the DHR source (which is Legacy), and the DPR source (for Dual-Pol).

Next, since the HPE and Bias HPE products are mosaics, they are identified as HPE and BHPE on the SCAN FFMP menu. Keep in mind there's a labeling error for these products. Both use Dual-Pol in their creation, however, their submenu reads "DHR MOSAIC". Just be aware that this is a typo.

Finally, the MRMS Radar Only mosaic is also available from the SCAN FFMP menu.

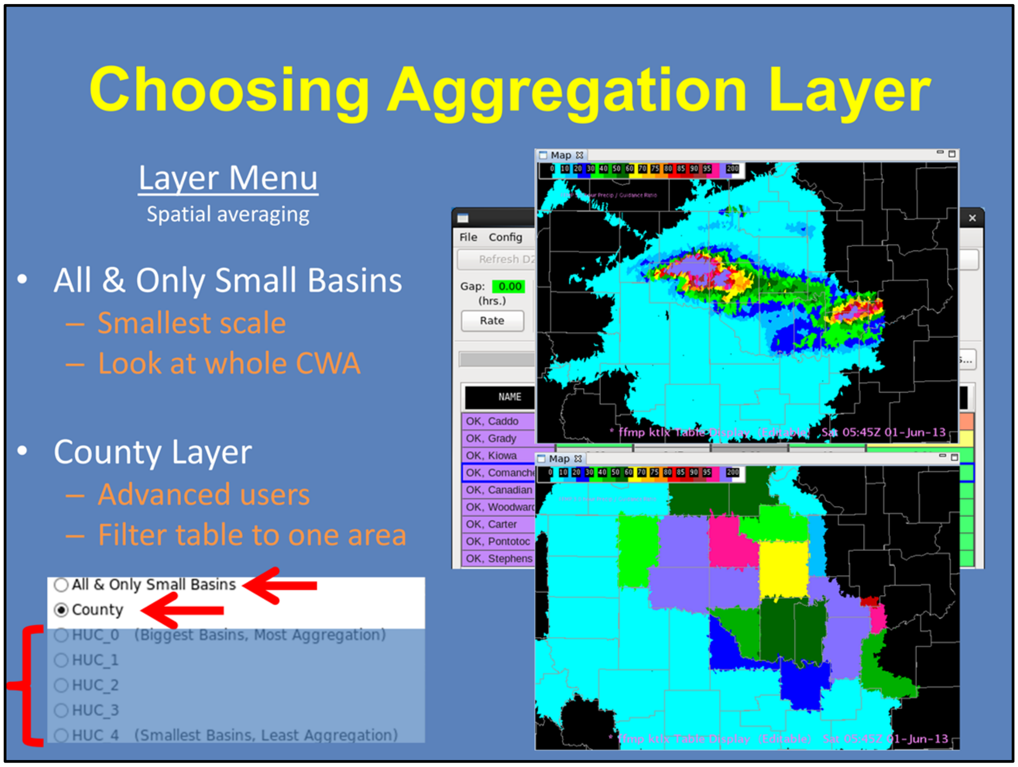

One of the most important steps in using FFMP is choosing the aggregation layer, which defines what spatial scale FFMP averages the QPE and FFG. The two layers we will focus on for flash flooding are the All & Only Small Basins and County.

When you first open FFMP and begin your flash flood interrogation, we recommend starting with the "All and Only Small Basins" layer option. This will give you a simple look over the whole CWA on the most relevant hydrologic scale to see what basins stand out in QPE, ratio, or diff.

As you become more advanced with FFMP, you can switch to using the "County" layer option to organize the basins in the FFMP table by county. This is done to make it easier to find particular basins and virtual gauge basins. There are a lot more settings to pay attention to when using county layer, though, so just be careful. The HUC layers are collections of small stream basins for larger scales and are not used frequently for flash flood decision making.

FFMP Recommended Settings

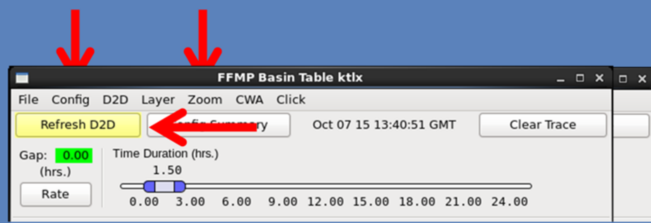

	All & Only Small Basins	County
Zoom Menu: Zooming options when aggregation layer is clicked		
Maintain Layer	OFF	OFF
Only Basins in Parent	OFF	OFF
Config Menu: display options		
Link to Frame	ON	ON
Worst Case for Aggregate	ON	ON
Auto-Refresh	ON	ON

There are several FFMP Table menu options that can enhance your D2D display, as well as your FFMP basin table display.

First, the Zoom menu controls how FFMP zooms into smaller basins when an aggregation layer is clicked in the table. In either “All and Only Small Basins” or “County” layer, we recommend these options be turned off. They do not have an effect on your display when using “All and Only Small Basins”. But if you use “County” layer, with the “Maintain Layer” option OFF, the D2D will not maintain the county layer and will instead show the individual basins. Setting “Only Basins in Parent” to OFF with the “County” layer allows any neighboring basins outside of the county to be displayed in D2D, so you can see flash flood threats crossing the county line.

Next, the Config menu helps with general display of the data. The “Link to Frame” ensures the D2D and the table are kept in sync when stepping through multiple frames.

Next is the “Worst Case for Aggregate” option. This option only comes into play when you have chosen a layer larger than “All and Only Small Basins”, like “County”. When turned on, this option sets the values in the FFMP table to show the “worst case” value for **any** basin within the larger aggregate layer. However, be aware that the “worst case” values may not always be within the same basins.

Finally, there is the Auto-Refresh option. This automatically updates the D2D display with any configuration changes made to the FFMP table. We recommend this be turned ON. However, if you notice performance problems with FFMP, you may consider turning it OFF. When it is off, you need to remember to click on the “Refresh D2D” button to manually update the display after making changes.

Back to Top

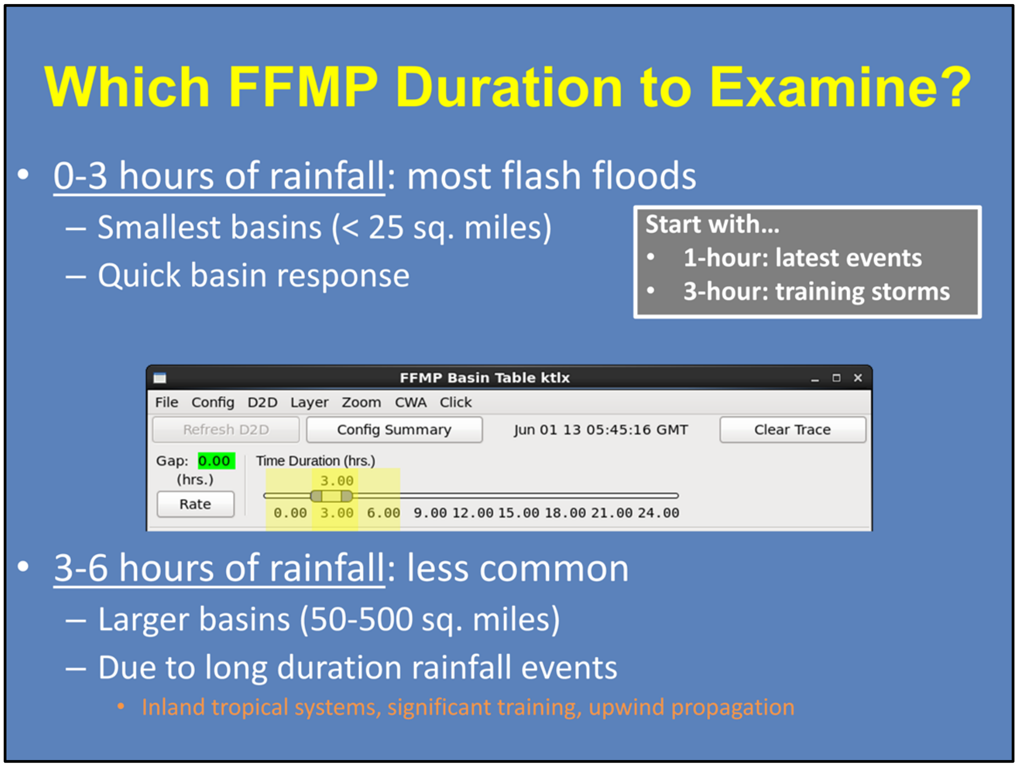

Now that we have reviewed the FFMP settings, let's talk about what durations are good to examine.

The majority of flash flood events take place due to less than 3 hours of rainfall, sometimes less than an hour of heavy rainfall. This is because they occur in basins smaller than 25 square miles. These tiny basins have quick responses to the rainfall, and thus inundate rapidly. Therefore, we suggest you focus on the 1-hour duration for the latest events, and the 3-hour duration for training storms. The duration slider bar feature in FFMP makes this easy to do.

However, certain meteorological environments are conducive to flash flooding larger basins, say 50-500 square miles, and thus require a longer duration to get things going. Inland tropical storms, significant cell training, and upwind propagation along a quasi-stationary boundary are examples of long duration heavy rainfall events that may result in flash flooding of large basins. In these types of set-ups, in addition to looking at 1 and 3 hour duration, it would be wise to also check out the 6-hour duration information from the basin table.

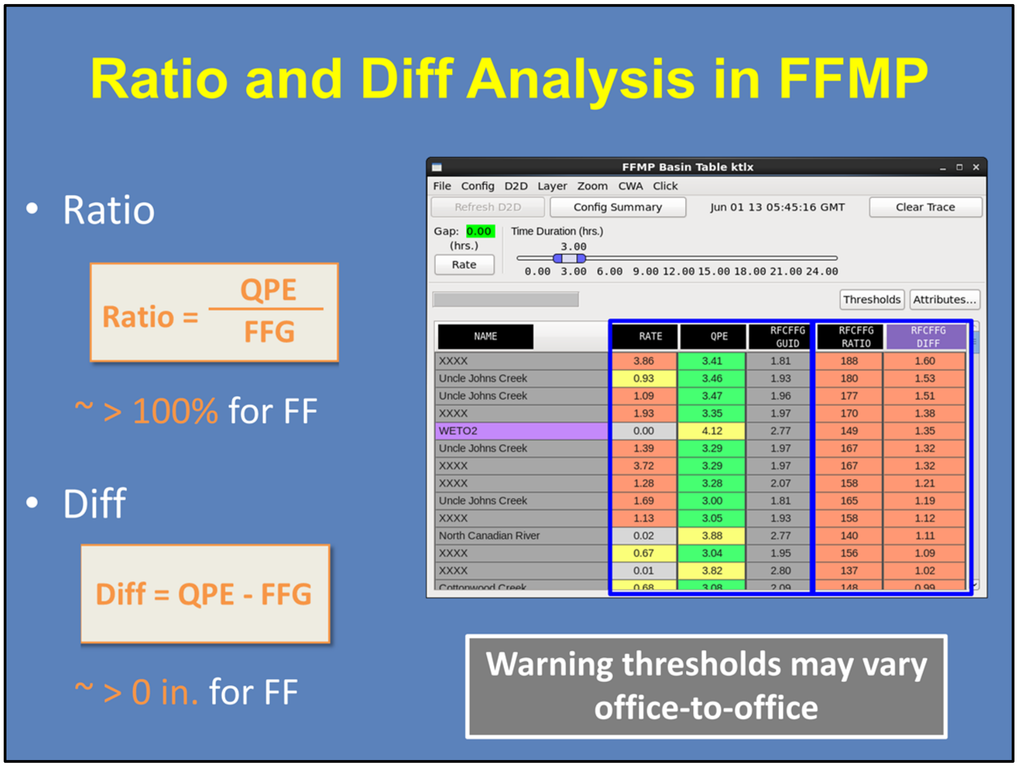

Okay, so you have your settings the way you want, and you know the duration you want to examine. So what should you look at? In addition to instantaneous rate, basin-averaged QPE, and basin-averaged FFG, FFMP has two other options for what can be displayed in the table: Ratio and Difference.

By default, these two take into account the Flash Flood Guidance values, and thus, are useful for analyzing exceedance threat. So let's take a look.

For Ratio, it is QPE divided by FFG. So, as Ratio approaches and exceeds 100%, that means QPE is near or exceeding FFG, and thus, the theoretical flash flood threat increases. To calculate the Difference, it is QPE minus FFG. So as Diff approaches zero or becomes positive, similarly, the theoretical flash flood threat increases.

But remember, all QPE sources have uncertainty and RFC FFG accuracy varies significantly across the country and over time. So you may find that ratio and difference warning thresholds vary from office-to-office. For instance, at some WFOs, flash flooding may typically start at 0.5 inches over FFG while another may start near flash flood guidance. But let's take a minute to review an example of QPE, ratio, and difference values.

Back to Top

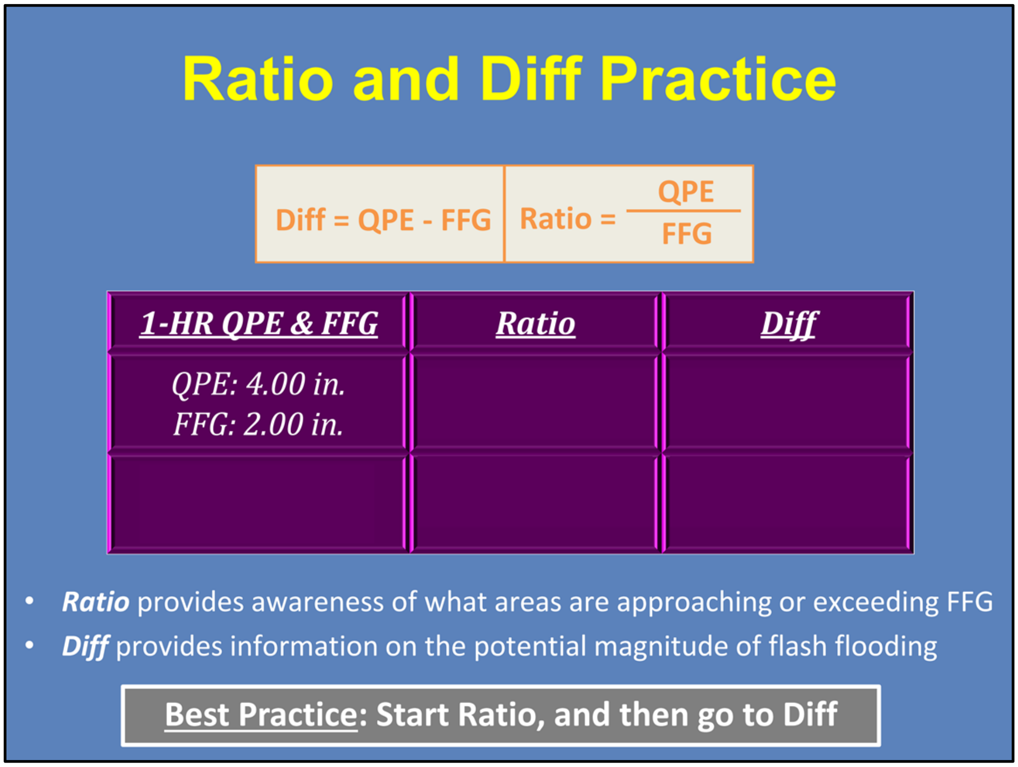

Alright, let's say we have a hypothetical basin that has basin-averaged rainfall of 4 inches in one hour, and the flash flood guidance is 2 inches in one hour. Thus, the Ratio would be 200%, and the Diff value would be 2 inches. Now imagine that for a different rainfall event, the same basin receives 1 inch of rain in an hour, and the FFG is only 0.50 inches.

The Ratio is still 200%, but the Difference is now 0.5 inches. This Ratio value could lead you to believe a significant flash flood was possible, as in the first case. However, comparing the two Difference values, the 1st event would have much more significant flash flooding given that FFG was exceeded by 2 inches, rather than only 0.50 inches during the second event.

Ratio can be used as a quick awareness tool for basins that are close to or already exceeded flash flood guidance. While, the Diff values give information on the potential magnitude of the flash flooding.

After identifying your areas of precip using QPE, we recommend that you start with viewing the Ratio, in order to pinpoint areas that may be approaching or exceeding FFG. Then, switch over to the Difference to tell how much you are over or under in those areas.

Back to Top

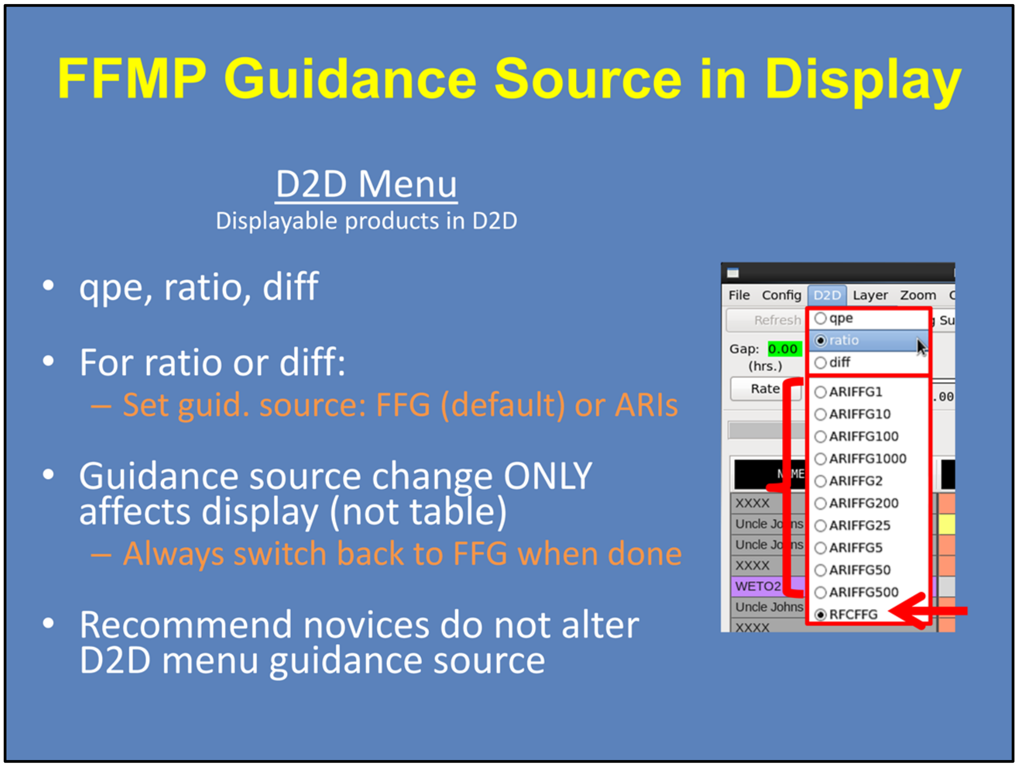

The last menu to discuss is the D2D menu, which determines what is being displayed in D2D. The three options are QPE, Ratio, and Diff, with the default being QPE.

Based on the best practice from the previous slide, it is usually good to start with QPE to get a feel for the high precip areas, and then move to Ratio and Diff to analyze the flash flood threat. When you switch the D2D menu option from QPE to Ratio or Diff, FFMP wants to determine what guidance source is being used for the ratio calculation. The default is RFC FFG, as shown here.

However, Average Recurrence Intervals (or ARIs) are a new guidance source option in FFMP. To force the D2D display to show ratio (or diff) calculated against ARI, simply choose one of the various ARI options. Keep in mind this change ONLY affects the display, and not the table values. Therefore, it can become confusing if you set the *display* to use ARIs to calculate ratio and diff, but your *table* uses FFG. So do NOT forget to always switch your D2D display back to FFG once you are done with the ARIs.

There is a lot to learn about ARIs before you start using them in AWIPS. For novices, we recommend you do not use ARIs in FFMP, and therefore do not alter the D2D menu guidance source. Rather, wait until you take WOC Flash Flood to learn more.

Back to Top

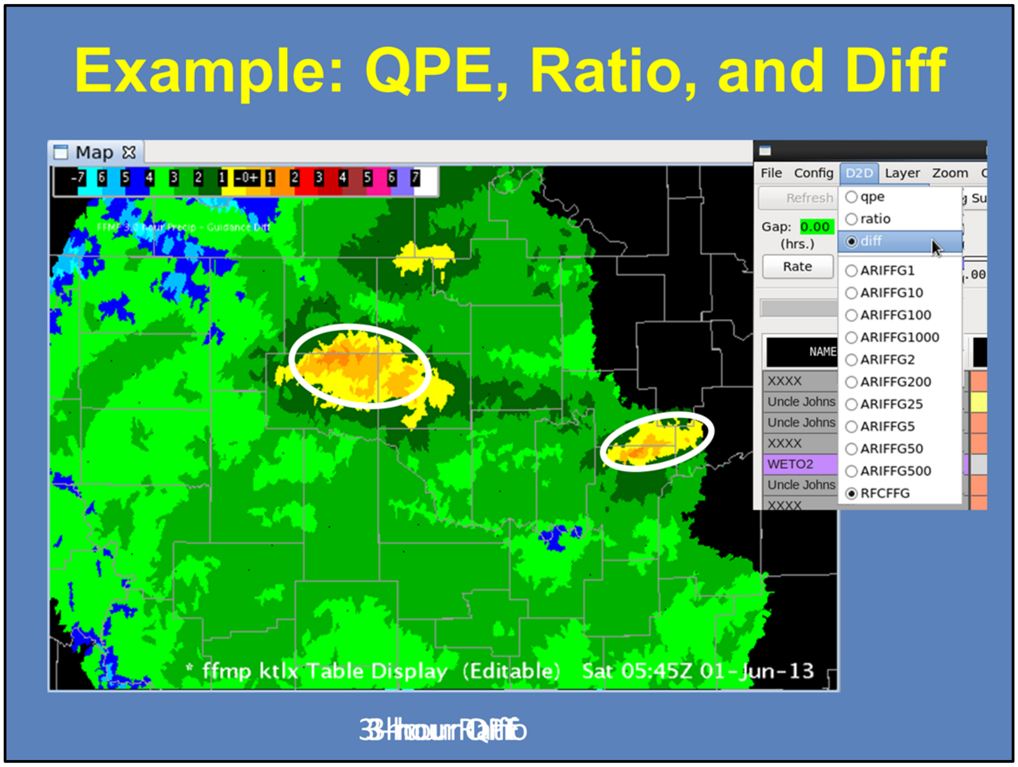

Here is an example of how to interpret the D2D options.

So FFMP defaults to displaying QPE. Simply looking at this output, we see there is a large area of greater than 1 inch in three hours, with isolated areas having upwards of 3 inches of rain in three hours. This information is useful for situational awareness, however it does not tell us anything about the hydrological response.

Therefore, your next move is to look at the Ratio product. Remember, for this, we are interested in areas that are approaching or have exceeded 100%. If rain is continuing in the area, then also consider the areas of 80 and 90%, since they are close to exceeding FFG. With this methodology, we have narrowed our threats to the circled regions.

Finally, use the Difference display to see by how much FFG has been exceeded. In this example, within our areas of interest, we have generally exceeded FFG within 1 inch. But there are some areas exceeding by 1-2 inches, which is where the more significant flash flooding threat is located.

So you see how this process helps you narrow down your flash flood threat, while providing details on magnitude that may be useful when considering your warning text.

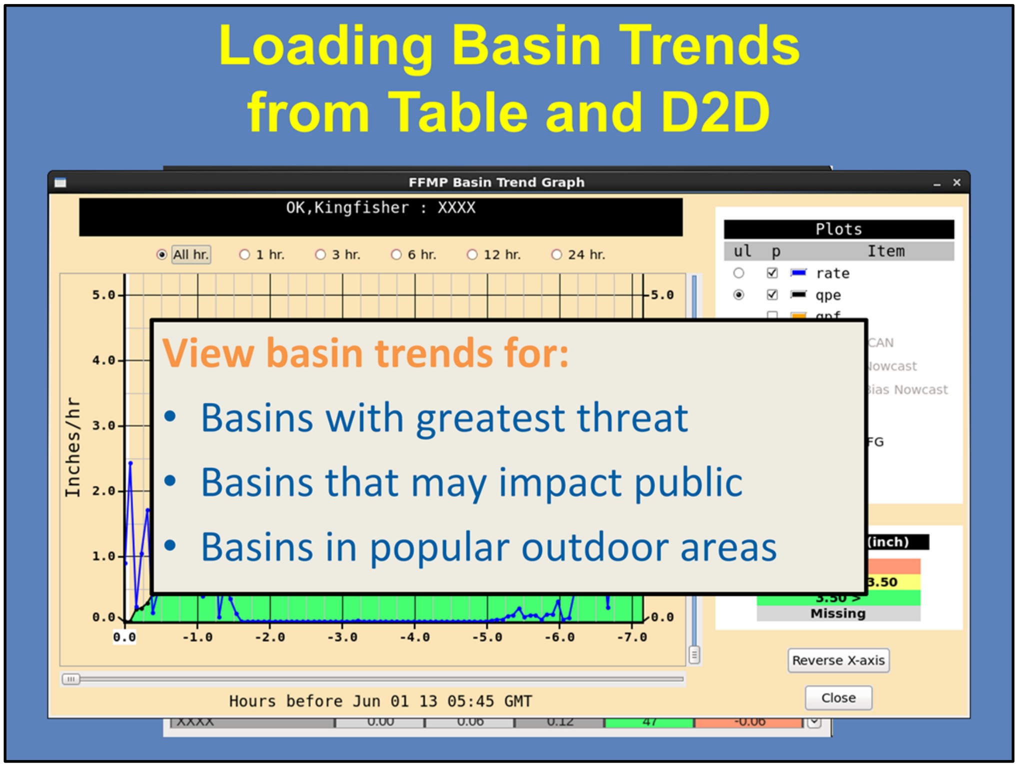

By this point you have zeroed in on the primary threat areas using Ratio and Diff and by monitoring rain rates. The next useful functionality in FFMP are basin trend graphs.

To load a Basin Trend, there are two options: First, you can load it by right-clicking on a small basin name from the basin table. Second, you can set the Click menu option to “Basin Trend”, then go to the D2D pane with FFMP and make the display “editable” (by middle clicking on the text in the legend), and then right-click on any basin in the display to load a basin trend for that basin.

Because there could be tens of thousands of small basins in your localization, it is best to focus on basins that: 1) have the greatest current or projected threat; 2) perhaps those basins that might significantly impact the general public (like urban basins); or 3) basins in a National Park that normally contain numerous hikers and campers.

Basin trend graphs are critical to interpreting information on the timing and relationship between the QPE and guidance for different durations. And with time, you will become more familiar with using them. We'll start you off with an example on the next slide.

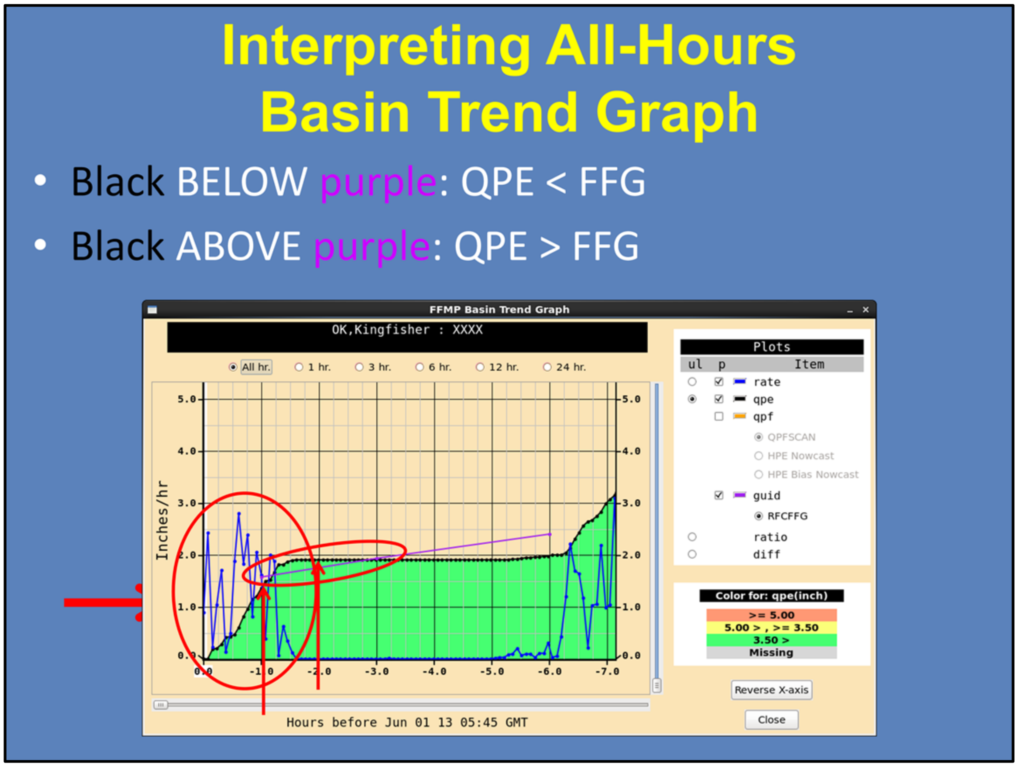

Okay, let's take a look at this basin, which is only three-hundredths of an inch away from exceeding the 3-hour FFG and is currently experiencing instantaneous rain rates of 0.89 in/hr. As a warning forecaster, I would like to know when within the three hours 1.90 inches of QPE fell, so I right-click on the basin to load a basin trend graph shown here.

First, you want to look at the blue line, which is the instantaneous rate trend. Each blue dot represents the instantaneous rate for a particular volume scan. From this we see that rates of ~ 2 in/hr occurred primarily over the last 1.5 hrs, and there was no precip 3 hrs ago and 4 hrs ago.

Next, the black line is the precip accumulation for different durations. You will notice the accumulations increase every time there is an instantaneous precipitation rate > 0. The instantaneous rate is multiplied by the volume scan time step in order to increase the accumulation. To interpret this line, we see about 1.3" has accumulated over the 1-hour duration, while 1.9" have accumulated over the most recent 2-hour duration. We see the 1.9" accumulation lasts through the 5-hour duration, because there was no precip falling between 2 and 5 hours ago.

Finally, there is the purple line, which shows FFG for the 1-, 3-, and 6-hour durations. Whenever the black QPE line is BELOW the purple FFG line, QPE is less than FFG for that duration interval. When the black line is ABOVE the purple line, QPE is greater than FFG. Here, QPE is always below FFG, except for durations between 1 and 3 hours where FFG is exceeded by about 0.25" for the 1.5-hour duration. This may be enough to cause flash flooding, particularly since the instantaneous rates are continuing at the current time, and the longer duration FFG values (like 3- and 6-hour) are going to be exceeded more and more as that continues.

Now we're going to take a few minutes and let you have some practice with basin trend graphs. The following quiz is NOT graded.

Back to Top

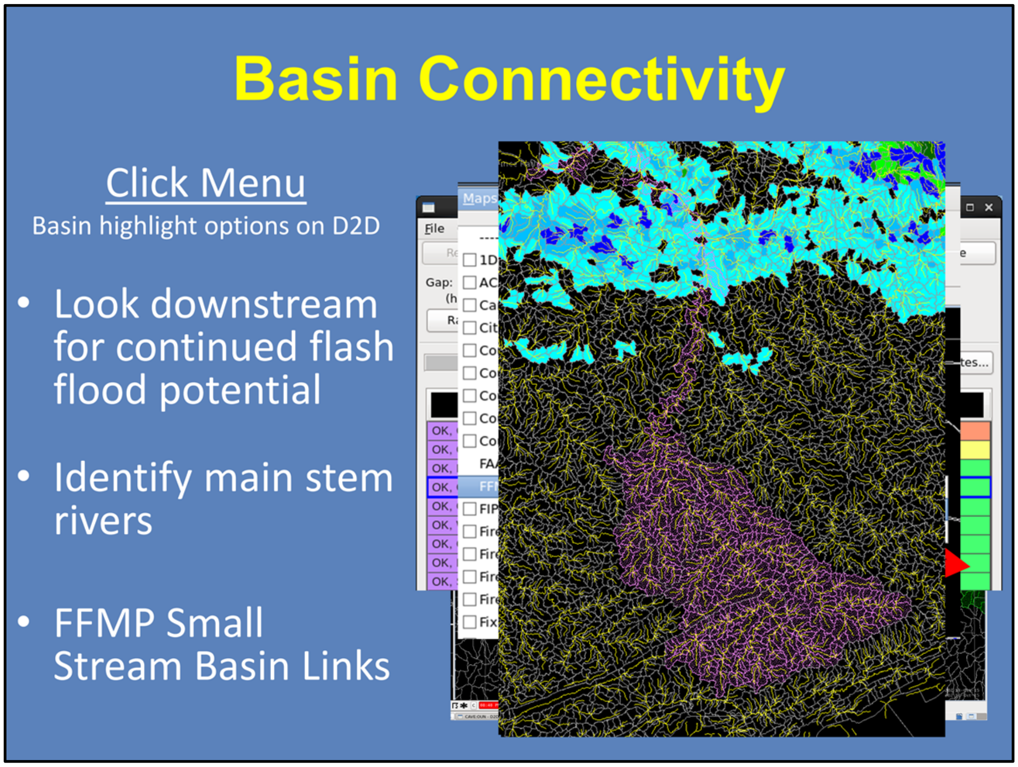

Finally, FFMP allows you to see basin connectivity on the D2D display as configured in the "Click" menu. Once you have pinpointed your current threat area, it is important to look downstream to see where the runoff will go. If the current accumulation is great enough and the downstream basins are flashy, those downstream basins can have flash flooding even without receiving a drop of rain. Knowing this can help you adjust your warning polygons to account for the future threat.

To do this, simply select "Downstream" from this menu, and then go to the D2D display and make the FFMP display editable. Once editable, your right-click button will highlight all basins downstream of the basin you selected. If it is hard to see the highlighting, you can change the color of the trace, like I did to green. And whenever you want to get rid of your trace or change the type of trace, simply "Clear Trace" on the FFMP Table.

Additionally, you may want to identify major main stem rivers. Since they typically don't *flash* flood, this may help you pull basins out of your warning. To do this, use the upstream and downstream option from the menu. Here is an example where the star denotes the selected basin. You can see the large area upstream of the basin that is feeding into that point. And then where it goes downstream to the north.

Finally, you can also visualize flow by overlaying the "FFMP Small Stream Basin Links" from the Map menu in D2D. I made them yellow in this graphic.

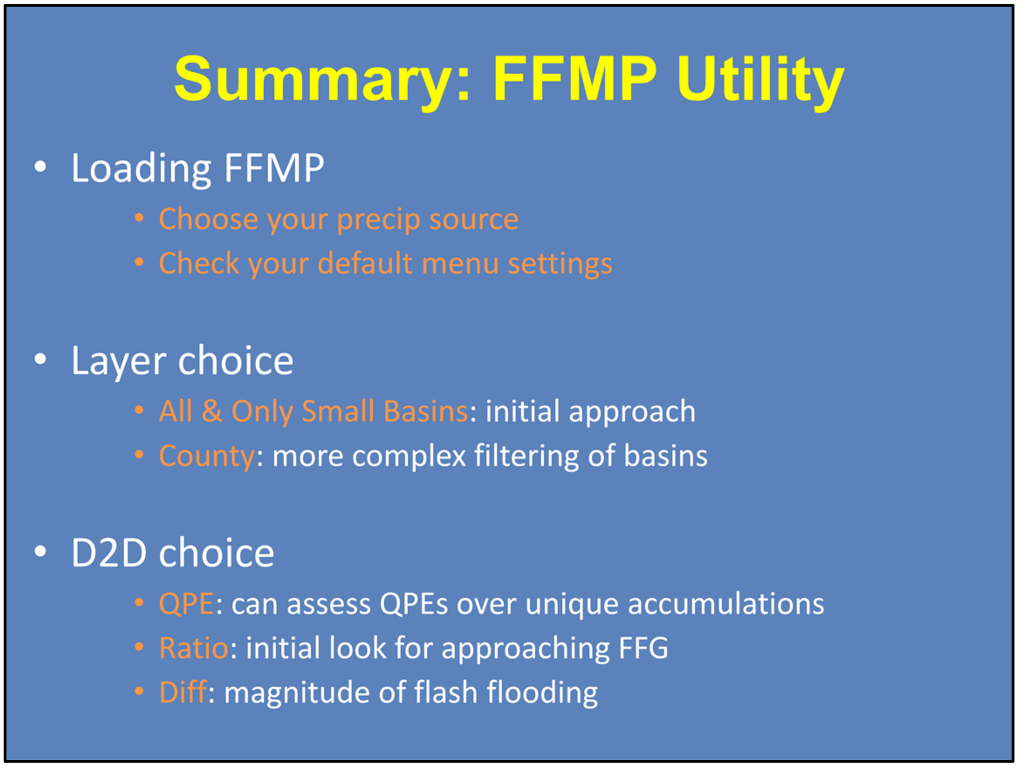

Because FFMP has a number of unique strengths, we focused this lesson on using it to its fullest in flash flood warning operations.

First, when loading FFMP, make sure to follow the guidance in the "Choosing Your Precip and Guidance Sources" lesson to consider all of your available precip sources. Also, make sure the menu settings across the top of the FFMP table are what you would like.

Next, start using the All & Only Small Basins layer to identify areas where QPE is approaching or exceeding FFG. You may change to County layer when you need more complex filtering of basins in the table.

Within FFMP, D2D can be configured to show any one of three options. QPE allows you to assess things like HPE, Bias HPE, and MRMS accumulation durations that aren’t usually readily available. The ratio product is one the best ways to identify areas of flash flooding threat so we recommend starting there, and using Diff to help assess the potential magnitude of flash flooding.

Back to Top

Summary: FFMP Tools

- Basin Trend Graph
  - See temporal trends: precip timing
  - Easy visualization: comparing QPE to FFG and gauges
- Basin connectivity
  - Downstream Basin Trace: where threat may evolve
  - Up/Downstream Basin Trace: identifying main stem rivers
  - Small Stream Basin Links: visualize flow

FFMP has a lot of useful functionality, as long as you know how to use it.

The Basin Trend Graph allows you to see temporal trends for rainfall rate, accumulation, and Flash Flood Guidance for a selected basin. As well as, provides easy visualization when comparing QPE to FFG, and to gauges when using VGBs.

FFMP also has basin connectivity features to help identify where the flash flood threat may be evolving, where main stem rivers exist versus headwaters, and how to visualize flow outside of FFMP.

This concludes this lesson. When you are ready, please move onto the next slide to take the quiz and receive credit on the LMS.

Back to Top

Hello, my name is Jill Hardy and welcome to this lesson on Flash Flood Warnings. Already in this course we have discussed the precipitation estimation products available, fundamentals in flash flood meteorology and hydrology, and how to utilize FFMP and the precip sources.

But before you arrive at the workshop, it is vital that you understand some flash flood warning fundamentals. This lesson will focus on the polygonology and warning text fundamentals, in order to give you the last piece of the puzzle for great flash flood decision making.

Back to Top

Learning Objectives

- By the end of this lesson, you will be able to:
  - Identify when to issue a Flash Flood Warning (FFW) versus Flood Advisory or Areal Flood Warning
  - Identify appropriate polygon sizes and FFW durations
  - Identify when to use the automated basin list in WarnGen
  - Identify appropriate basis and call-to-action details to include in a FFW
  - Identify how to follow-up a FFW
  - Identify the criteria for a Flash Flood Emergency

Here are the learning objectives for this lesson. When you have finished reading them, please continue to the next slide.

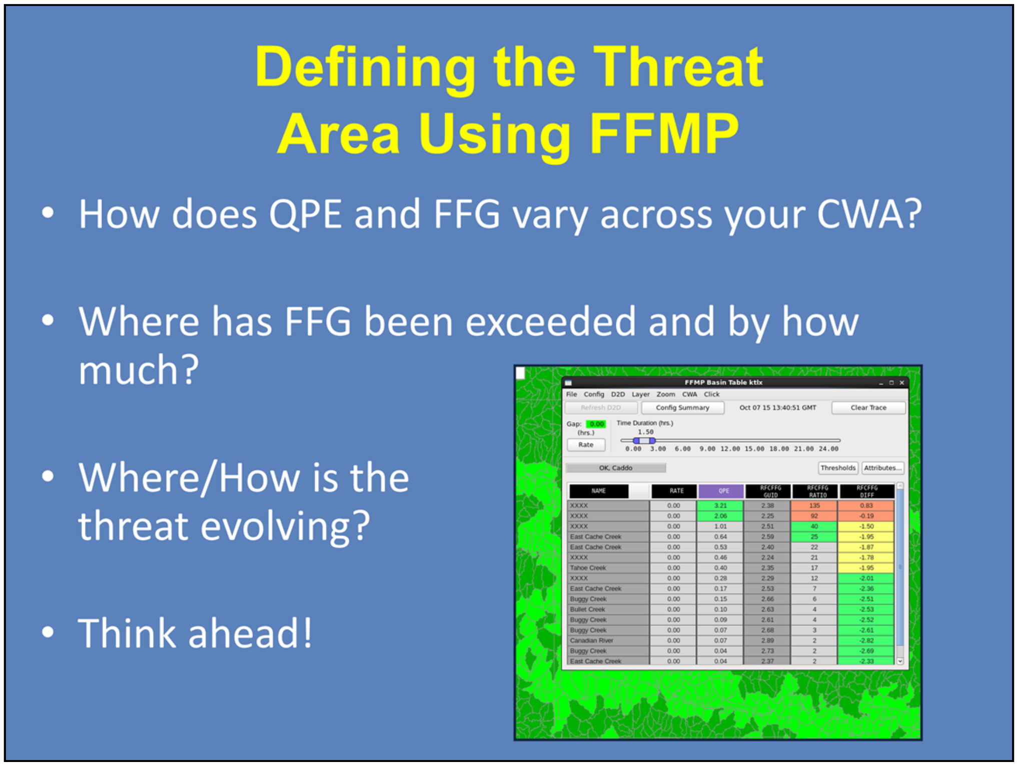

Before we jump into the details of creating polygons and text warnings, there are several questions you should always be asking yourself when defining the threat area.

First, how does your selected QPE source and flash flood guidance vary across the CWA? Is there an area of very high QPEs moving into an area of low FFG? Or, where has flash flood guidance been exceeded and by how much? Do you have an ongoing event? Next, where and how is the threat evolving? Again, is the event moving into an area of low FFG? Or, are training storms a concern?

The main take away is to always be thinking ahead! FFMP focuses heavily on what is happening now, based mainly on QPE. But as a forecaster, you must mentally extrapolate storm movement and threat evolution to generate proper lead time, particularly for rapid runoff in urban areas.

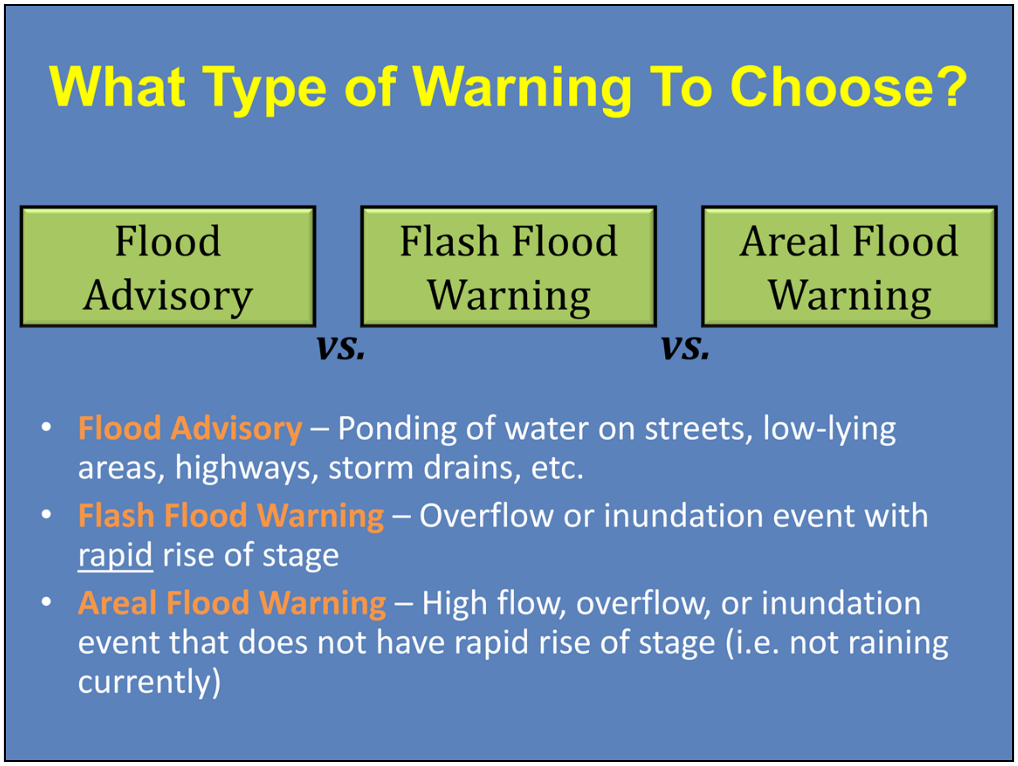

So let's start moving through the things to think about while on the hydro desk. Let's say you've diagnosed that the potential for flooding is likely…So what do you issue?

There are a few routes you can take: a Flood Advisory, Flash Flood Warning, or Areal Flood Warning. Make sure to talk to your office to see if they have a protocol in place.

Generally, a Flood Advisory covers any sort of ponding that is not life-threatening. A FFW should be used when there is a RAPID rise of water, within 6 hours. Whereas, an Areal Flood Warning would be used if there is high flow, but it is not a rapid rise.

So remember, for a FFW, you're looking for a rapid rise of swift-moving water. If this criterion is not met, one of the other two is probably a better choice.

A best practice: Let's say you are expecting a widespread, long duration rainfall event with marginal rain rates. Putting out a 6 to 24-hour Flood Advisory or Areal Flood Warning may be an effective product. But let's say that during this long duration event there are small areas of localized heavy rain rates that could lead to life-threatening flash flooding. Here, you would want to embed FFWs already under the Flood Advisory or Areal Flood Warning. This set-up properly explains the different threats. Once the flash flooding warnings have expired, Flood Advisories or Areal Flood Warnings can be continued for basins that are still seeing non-life-threatening, general flooding.

Back to Top

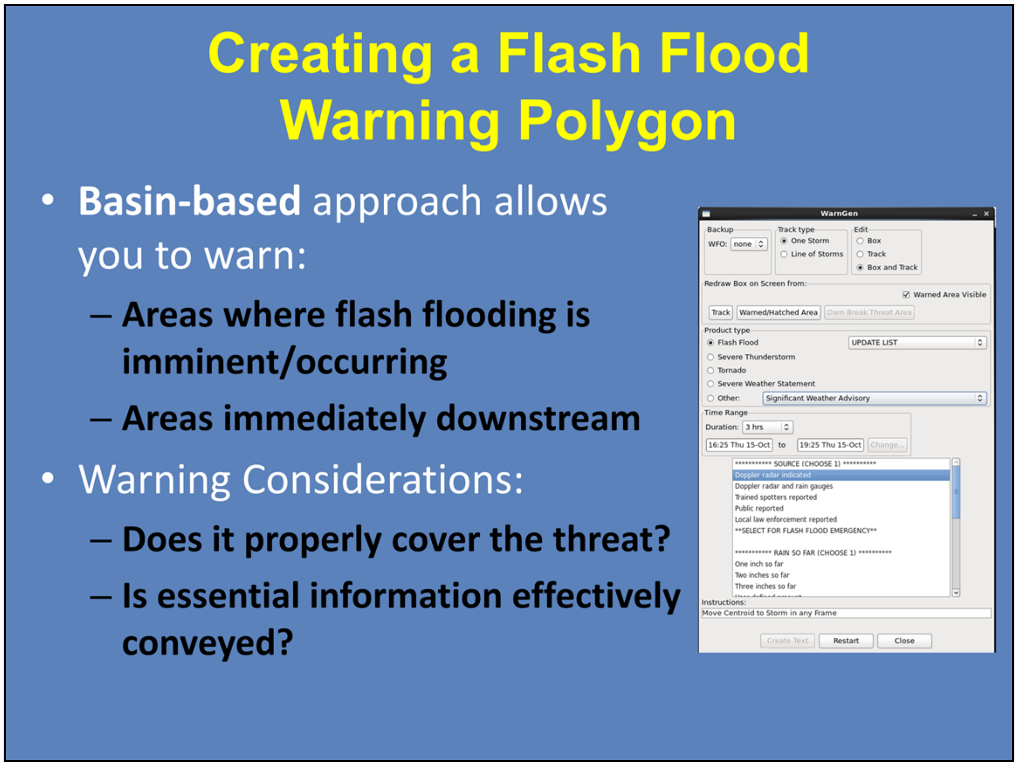

Okay, so you’ve decided that a FFW is necessary. So let’s look into some of the warning polygon fundamentals.

First, your warnings should be basin-based, rather than storm-based. Remember, flash flooding is a two-headed beast…You must consider the meteorological **and** hydrological factors. An area with heavy rainfall will not produce flash flooding unless the hydrological criteria are also met. Basin-based warnings allow you to warn the areas where flash flooding is imminent or already occurring, as well as areas immediately downstream.

Additionally, you should consider if the polygon properly covers the threat area…Not only right now, but in a few hours too. Where is it moving? What is the hydrology like there?

And, is essential information effectively conveyed? You don’t want the threat to be overshadowed by wordy warnings listing obscure basin names.

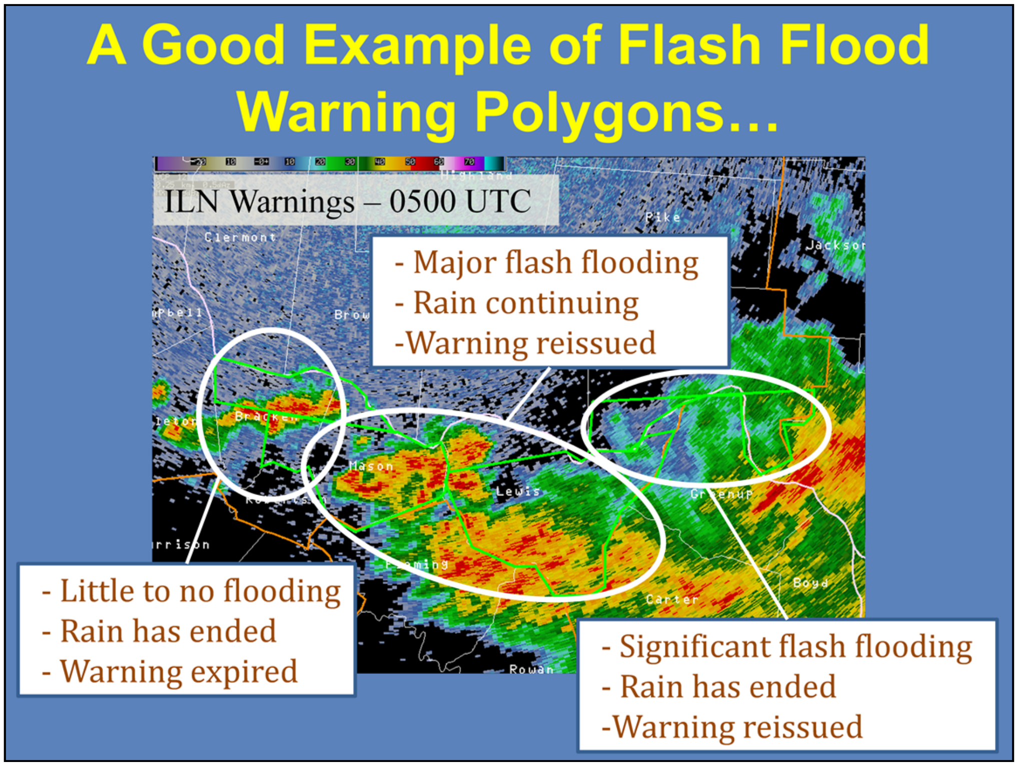

So let's look at a good example of flash flood warning polygons from the Wilmington, OH office. We start at 0500 UTC, and we see 4 active warnings. Let's go ahead and circle the different threat areas, combining the two warnings in the middle.

Now, let's move forward an hour to 0600 UTC. For the far western threat area, there was little to no flooding, and the rain has ended. So this warning was allowed to expire. Moving east, there was major flash flooding occurring at this time, and the rain is continuing. Therefore, the office reissued the warning, and combined the threat area. Finally, the far eastern area had significant flash flooding, even though the rain had ended. Therefore, the warning was reissued.

In this case, the warnings were properly itemized, so to explain the evolving threats. And notice how the 4 original warnings were made to expire at the same time, which helped in the reissuing process at the later time.

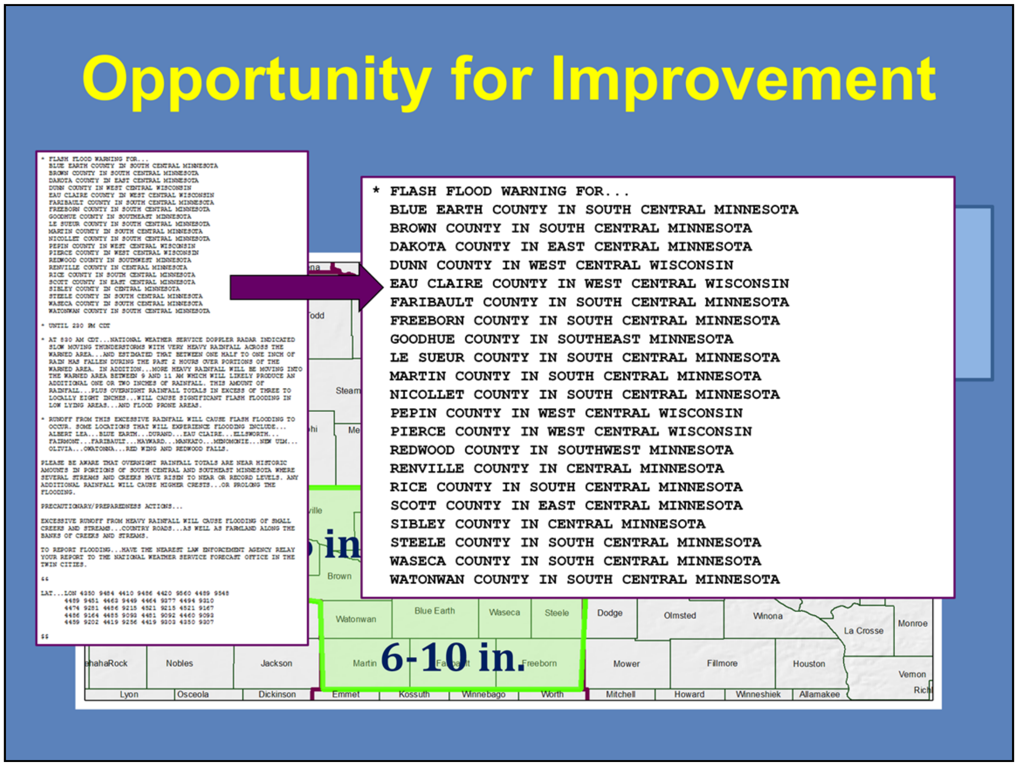

What about this example? Does the warning properly cover the threat area? And is the warning text effective?

In this case, there are 21 counties included in one warning. Would you expect the threat to be the same in all locations? Likely not.

In fact, there was a spread of 1-10 inches across the warning. Yet, all areas are receiving the same message. This is not ideal.

And in the warning text, all 21 counties are listed, which would be difficult for the existing systems to disseminate, like NOAA Weather Radio. Or even for the public to have to decipher themselves.

Back to Top

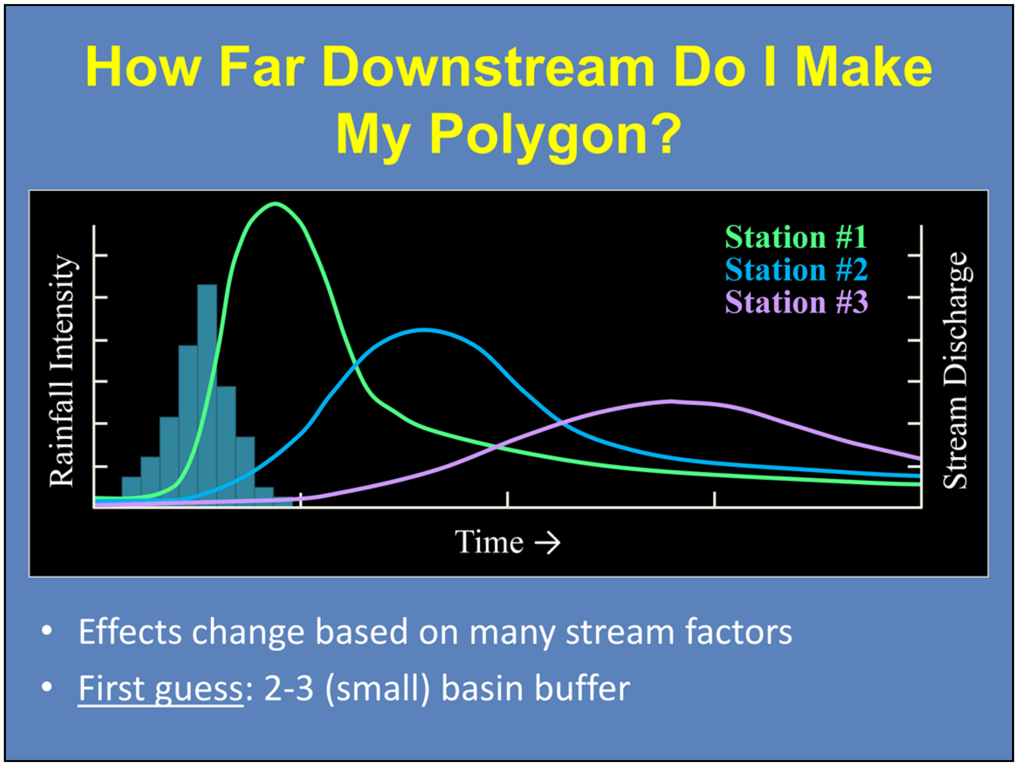

So you figured out your immediate threat area, but how far downstream is necessary to cover the evolving threat?

Here we have a plot of rainfall intensity over time. Let's say Station #1 is very close to where the rainfall fell. It is going to respond quickly to the rainfall, with a sharp jump in stage. Let's say Station #2 is a little further downstream, so its response is later than the first, with a more gradual rise in stage that isn't as high as Station 1. Finally, Station #3 is the farthest downstream, so it only sees a slight rise in stage.

This progression can be expected in most cases, however effects can change based on many stream factors. A good place to start in the absence of any local hydrological knowledge is to expand your FFW 2-3 basins downstream to account for runoff, not 2-3 counties! This is in addition to the expanding threat due to training storms and the short-term movement of precipitation areas.

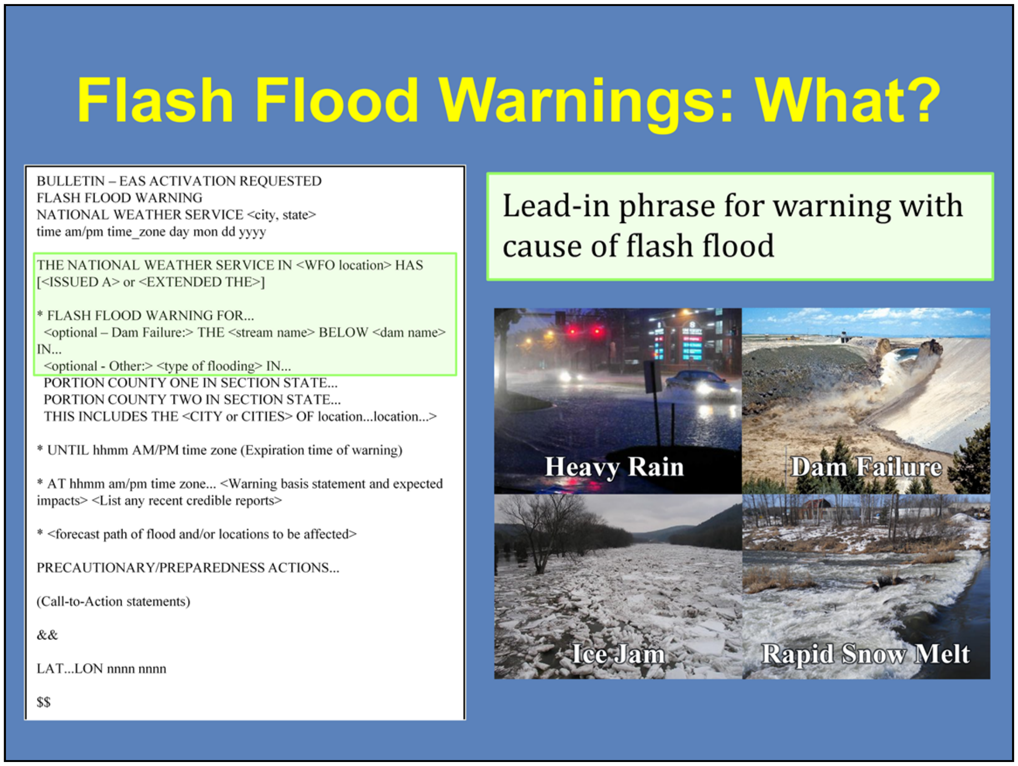

Okay, so you've drawn the perfect polygon, now let's start looking at the warning text best practices.

The first question is: What is the cause of the flash flooding? This will be the "lead-in" phrase for your warning.

In this course, you'll just focus on heavy rain. However, be aware that there are several other causes that you should learn more about from your office.

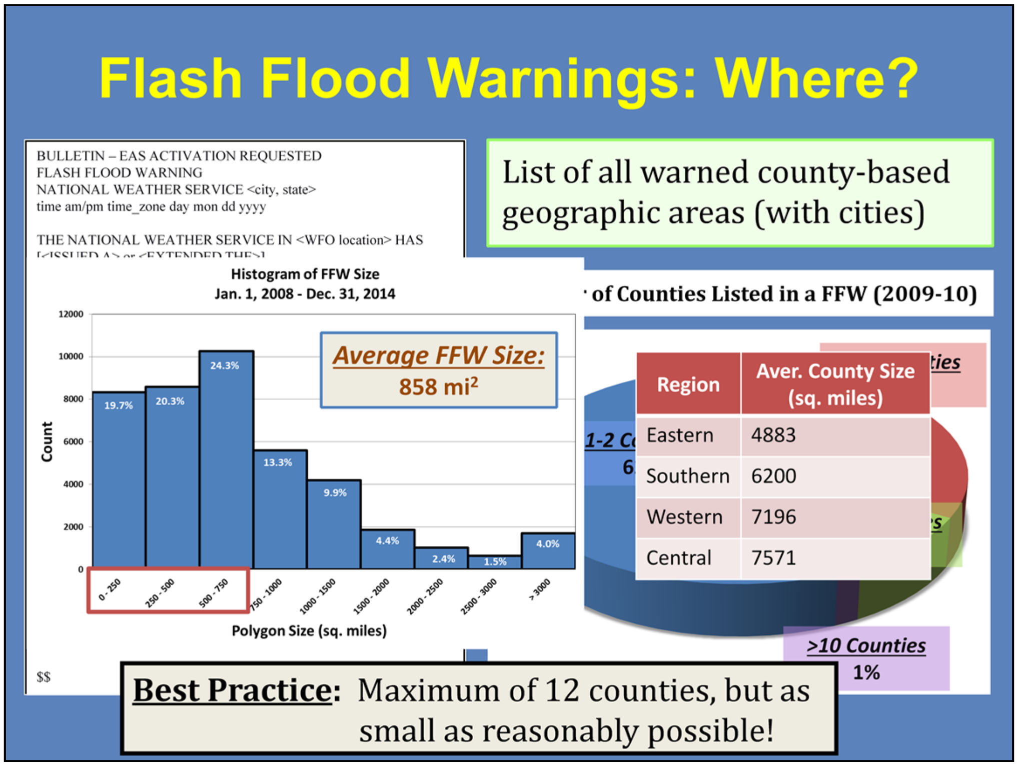

Then you have to add in the "where". Where is this threat imminent or already occurring?

When it comes to how many counties you should include, consider this pie chart. It is the number of counties listed in a FFW, using data from all FFWs for 2009 and 2010. This chart shows that the vast majority of FFWs include less than 6 counties.

Since county size varies a lot across the country, consider this histogram, which uses data for all FFWs for 2008 through 2014. It shows the distribution of warnings based on size, in square miles. Here you can see that the majority of FFWs are less than 750 square miles, with an average warning size of around 850 square miles. To put this in perspective, here is the average county size for each region in the CONUS. So the good news is that our warnings, on average, are quite a bit smaller than the average county. However, there were several warnings between 18 and 24,000 square miles. So we still need to minimize warnings of these sizes, due to the reasons described earlier: what's the likelihood of everyone in these warnings having similar threats?

To give a number, best practice is no more than 12 counties included in your warning, again, understanding this will vary with county size. But ideally, your goal as a forecaster issuing FFWs is to make the warning as small as reasonably possible to cover the threat area and how it will evolve with time.

Back to Top

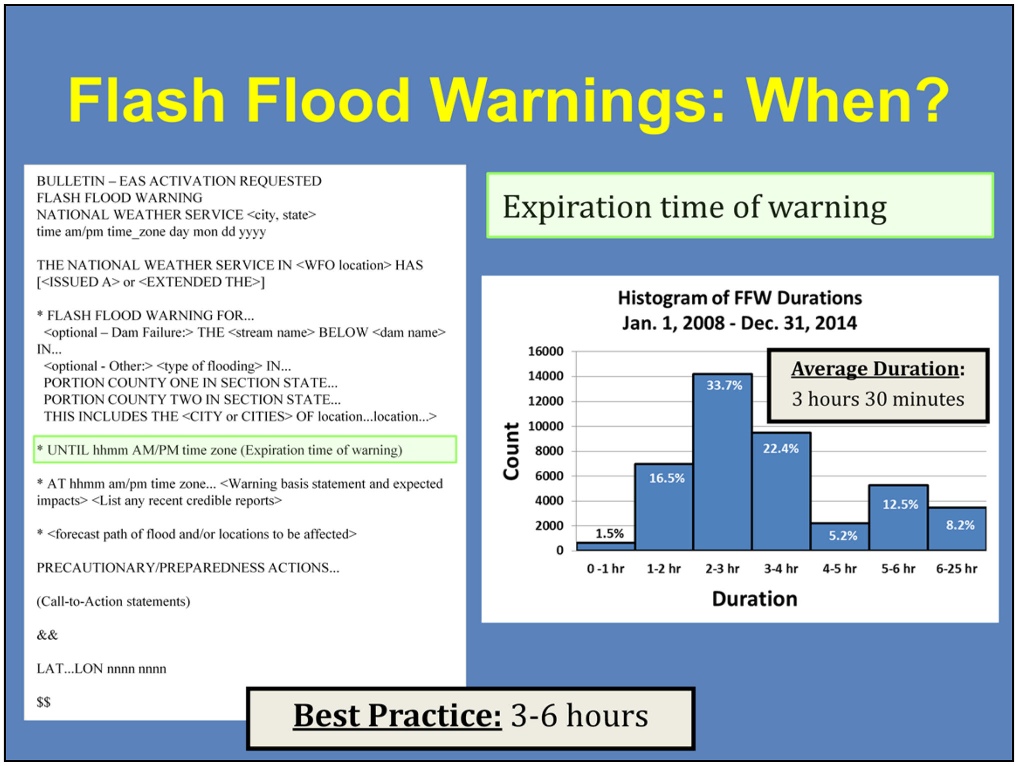

Next you have to add the "when". When should the warning expire?

This is another selection that will vary significantly based on the cause, location, if storms are training, etc.

Using data for all FFWs from 2008 through 2014, we see the distribution of warning durations. The average for this 7-year dataset was 3 hours and 30 minutes.

In the absence of unusual circumstances, a best practice is for FFWs to be between 3 and 6 hours. For routine FFWs, 3 hours allows for one hour for the event to begin and the rain to fall, one hour for runoff and the stage to crest, and one hour for the flood to recede.

But let's see what the directive says…

Back to Top

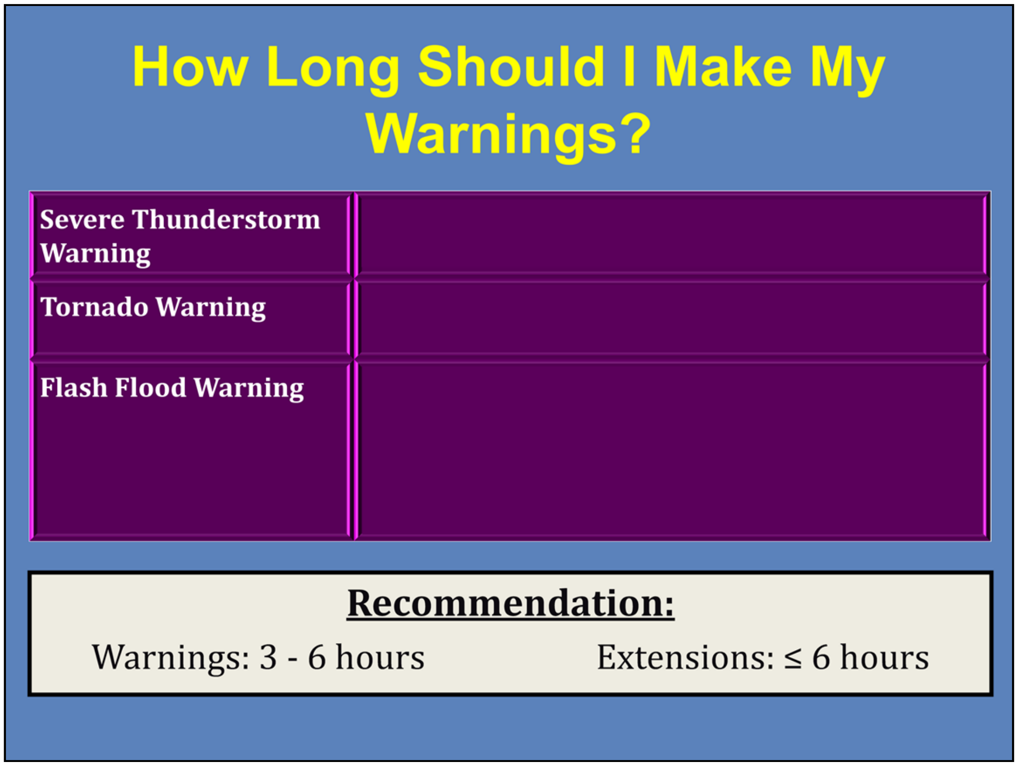

Pop quiz! Let's quickly review the different directives.

For a severe thunderstorm warning, what do you think the directive says? The answer is 30-60 minutes. How about a tornado warning? Know that one? It's 15-45 minutes.

Alright, now a FFW. Got a guess? Here is the answer: A flash flood warning will be valid from the time of issuance until the time when flooding (requiring immediate actions to protect life and property) is expected to end.

Hmmmm, so what does that mean? What would you say is a minimum? Maximum? Well, I can't give you an exact rule to live by, but here is what we recommend: For heavy rain threats, make your initial warning a minimum of 3 hours, as explained on the previous slide. The recommended maximum is 6 hours, if you expect repeated cores of heavy rain to move through the area.

In rare cases where there is long-term excessive rainfall where life-threatening flash flooding continues beyond 6 hours, extensions of no more than 6 more hours can be issued, if needed. Any longer than that, and you're getting into the realm of areal flood warnings.

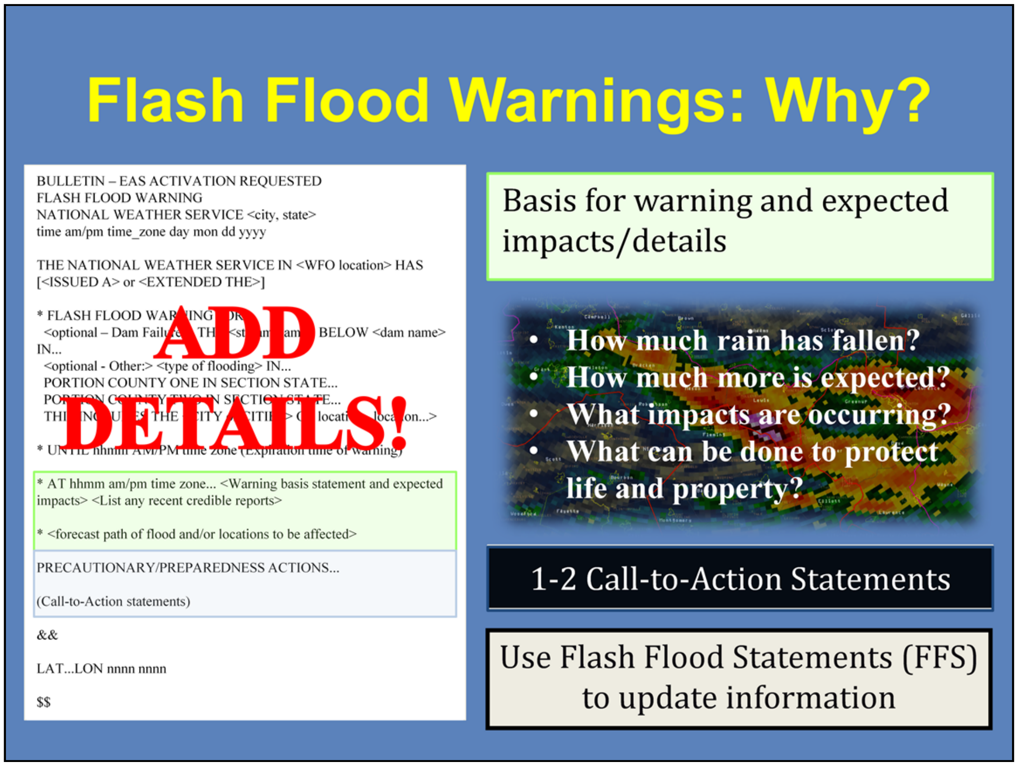

Finally, we need to explain the "why".

At this point, you may be thinking "Man, there's a lot to remember to forecast and warn on flash flooding!" But you know one good thing we've got going for us? A little bit of extra time!

Be sure to put in the details in the basis and call-to-action statements!

Answer these questions: How much rain has already fallen? How much more is expected? What impacts are occurring or can be expected? Include any relevant reports. And how can you protect yourself and your property?

Talk to your office about how many call-to-actions are good, but generally 1-2 per warning is best practice. You can include more if it's a significant event, long-lasting, in a metropolitan area, etc. Just do whatever properly disseminates the threat information.

Finally, don't forget to use Flash Flood Statements to update an ongoing event. Flash flood warnings are relatively long, recommended at 3 to 6 hours. Flash Flood Statements can be very useful in disseminating new information as the threat evolves.

Back to Top

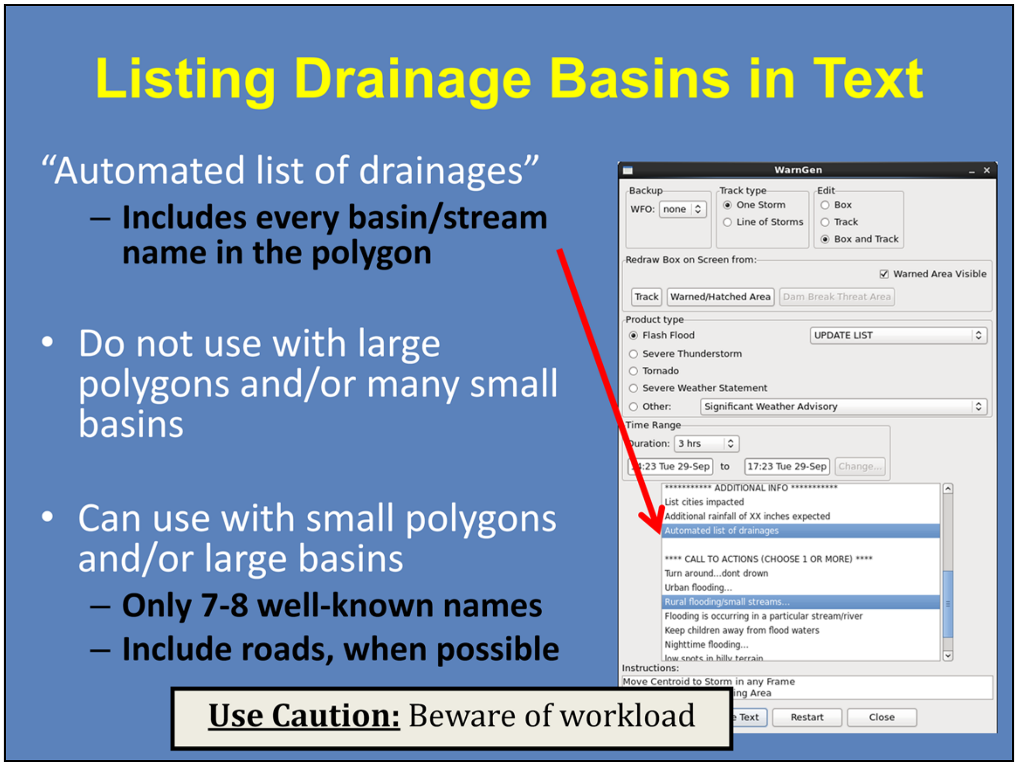

An important side note about one particular call-to-action.

Basin names can be automatically inserted into FFWs by clicking the "Automated list of drainages" option under the Calls to Action, shown here. What this does is include every single basin/stream name that falls within the warning polygon. There are positives and negatives to this.

For one, if a polygon is too big or in an area with small basins, this option can lead to a list of hundreds of basins. Since there is no geographic organization to this list, many of the basin and stream names may be unfamiliar to the general public. This can lead to un-needed and unwanted text in your warning. Reducing this list would require massive amounts of text editing, and simply not feasible during warning operations.

However, if a polygon is small or basins are rather large, this option may prove useful. Only keep drainage names of the creeks under the biggest threat and those well-known to the average customer, and/or those creeks and rivers that are well identified by signage for travelers to the area, if that information is indeed known. Try to reduce the number of names down to about 7-8. Also, include known road crossings that may be affected by flash flooding since the general public and media would recognize those even better than most creek names.

Simply use caution when considering this option. Determine if your hydrologic knowledge and workload can handle the edits needed to make the information pertinent to the public.

The "Flash Flood Emergency" Statement

- Criteria for using "Flash Flood Emergency"
  - **Short duration (1-6 hrs) precipitation causing major flash flooding over/upstream of populated area**
  - **Multiple swift water rescues**
  - **Stream gage data to major or rare levels**
  - **Total failure of a major high hazard dam**


Courtesy of Jeff Brown, The Republican (Mass.)

Other considerations: current road impacts, deaths, soil moisture

The last question is "When should I use a Flash Flood Emergency" statement? Well, to put it in perspective…Out of about 6000 FFWs per year, only about 5-10 are emergencies. Here are the criteria for a FF Emergency.

Basically, there must be an imminent or ongoing **elevated** threat to human life and/or **catastrophic** damage to property. Other considerations include road impacts, reported deaths, and soil moisture.

Remember: This is not a forecast! It should be issued only after you have reports. However, keep in mind that while you want to wait until you know the event is worthy of an Emergency tag, you want to declare an Emergency while it is still early enough in the event to be useful in terms of life and property impacts.

Back to Top

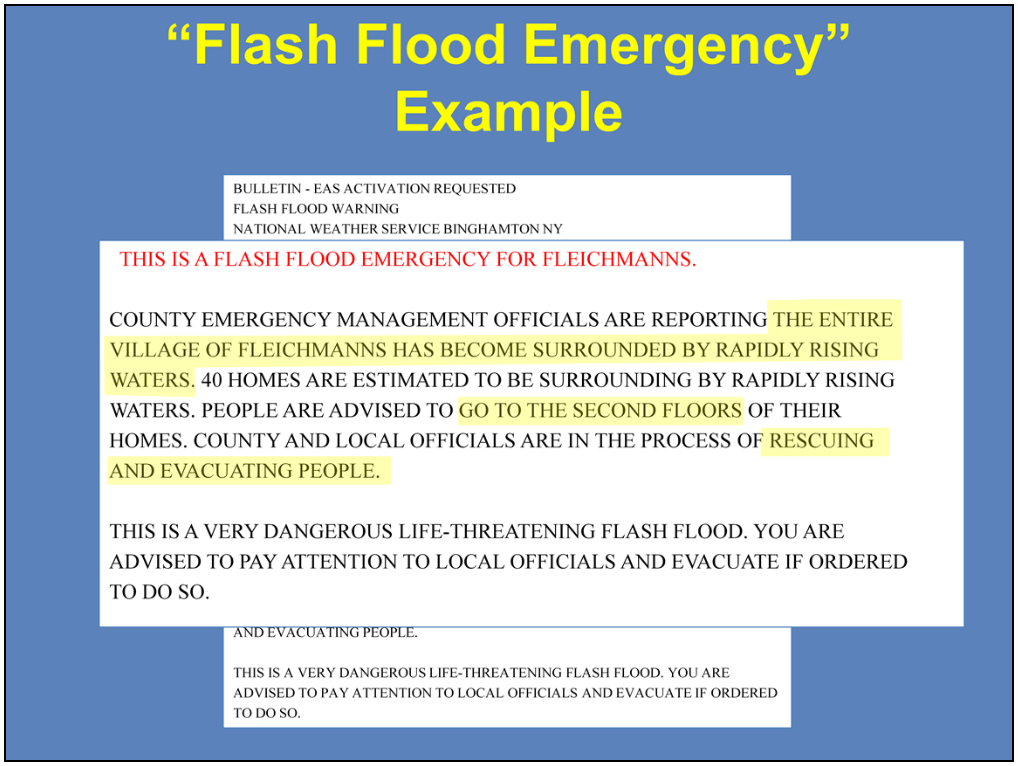

Here is an example of a Flash Flood Emergency statement.

Notice the language used: "…the entire village…has become surrounded by rapidly rising waters", "go to the second floors", "rescuing and evacuating people".

In these circumstances, if you are going to use the Emergency statement, you want people to change their behavior by using strong wording.

More information about Flash Flood Emergencies is covered in the WOC Flash Flood Course.

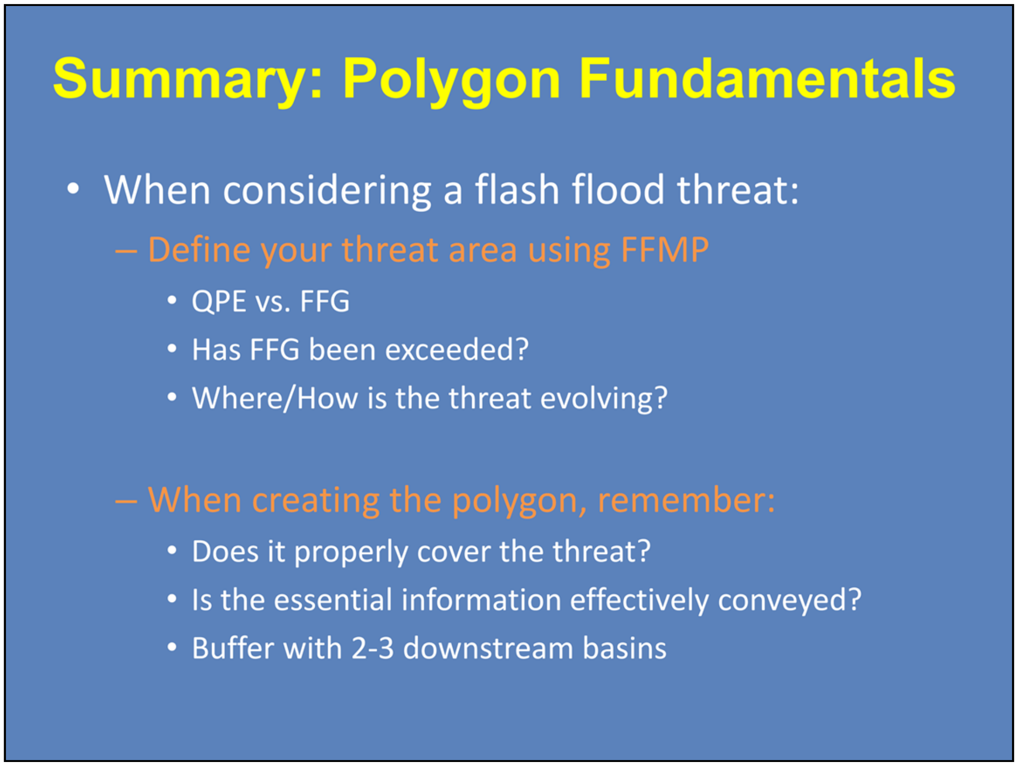

In summary, there are several things to think about at the start of your flash flood threat assessment.

Before even drawing your polygon, define the threat area using FFMP. Consider how QPE and FFG compare. See if FFG has been exceeded and by how much. Then, think about where the threat is evolving. Does the current storm motion, rain rates, and FFG values help point out the next area to expect impacts?

Once you start creating your polygon, ask yourself two important questions: Does the polygon properly cover the threat? And is the essential information effectively conveyed? Too large of polygons can lead to problems with impacts not being relayed to the correct audience, as well as cause alert systems to have issues reading the lists of county and basin names. Oppositely, too small of polygons can accidentally leave threatened areas out. Remember to account for routing effects by including a 2-3 basin downstream buffer to your polygon, in addition to the expanding threat.

Back to Top

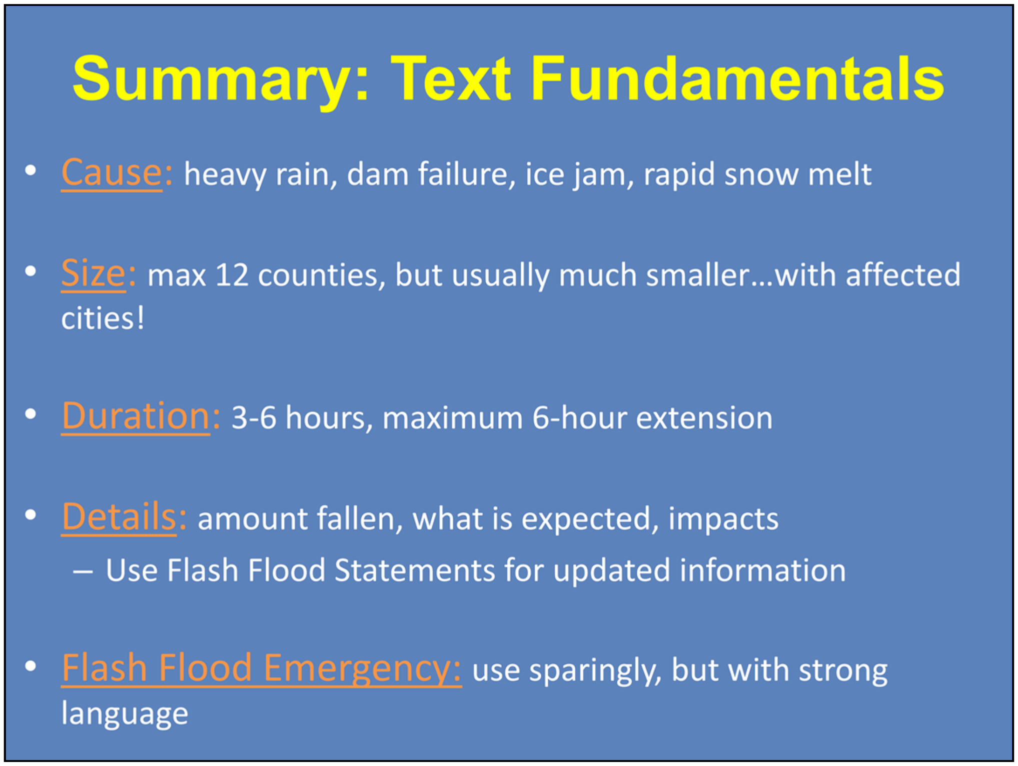

Once you are happy with your polygon, it's time to start thinking about the warning text.

1. Include the cause of the flash flooding, which is generally one of these four things. Talk to your office about how they handle each of these options.
2. As for size, you don't want any more than 12 counties listed, as it can make alerting difficult. But of course, try to make the warnings as small as reasonably possible, so that you are properly covering the threat. And, make sure to include any affected cities, since this is important, easily understood information.
3. FFWs, generally, should not be any shorter than 3 hours or any longer than 6 hours, and should not be extended more than 6 hours at a time. The event needs enough time to evolve, but also shouldn't be so long that the impacts are more areal flood-related.
4. Finally, and perhaps most importantly, make sure to take your time and include the necessary details about the event. Include how much rainfall has already fallen and what else is expected over the warning duration. And include current and expected impacts. Follow the warning with at least one Flash Flood Statement, including any reports and relevant rainfall and hydrological updates.
5. And hopefully you won't have to worry about it, but use Flash Flood Emergency wording sparingly. If you do meet the criteria and decide to use it, strong language should be used to notify the public of the elevated risk and to take action to protect themselves and their property.

This lesson was meant to provide useful recommendations when you are creating flash flood warning polygons and text. But don't forget to work with your local office to learn more about their protocol, as well.

Back to Top

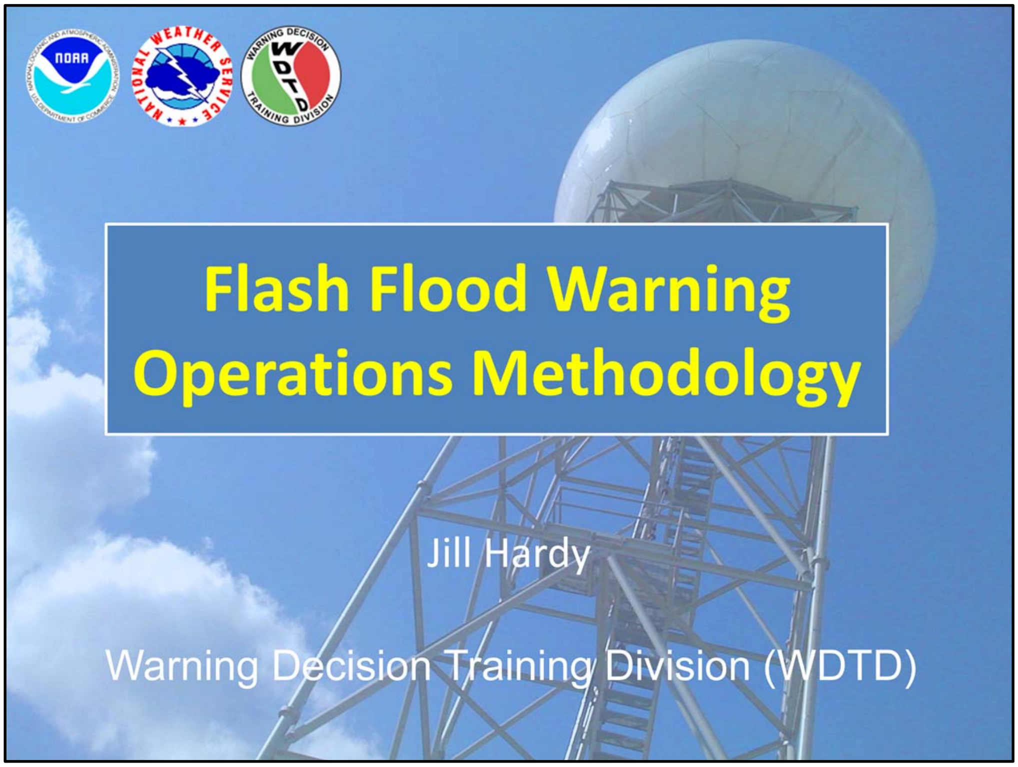

Hi, my name is Jill Hardy. This lesson is a brief summary of WDTD's recommended flash flood warning operations methodology. Basically, when you're the hydro warning forecaster on shift, what are the general steps and best practices to help you effectively issue flash flood warnings. Let's get started!

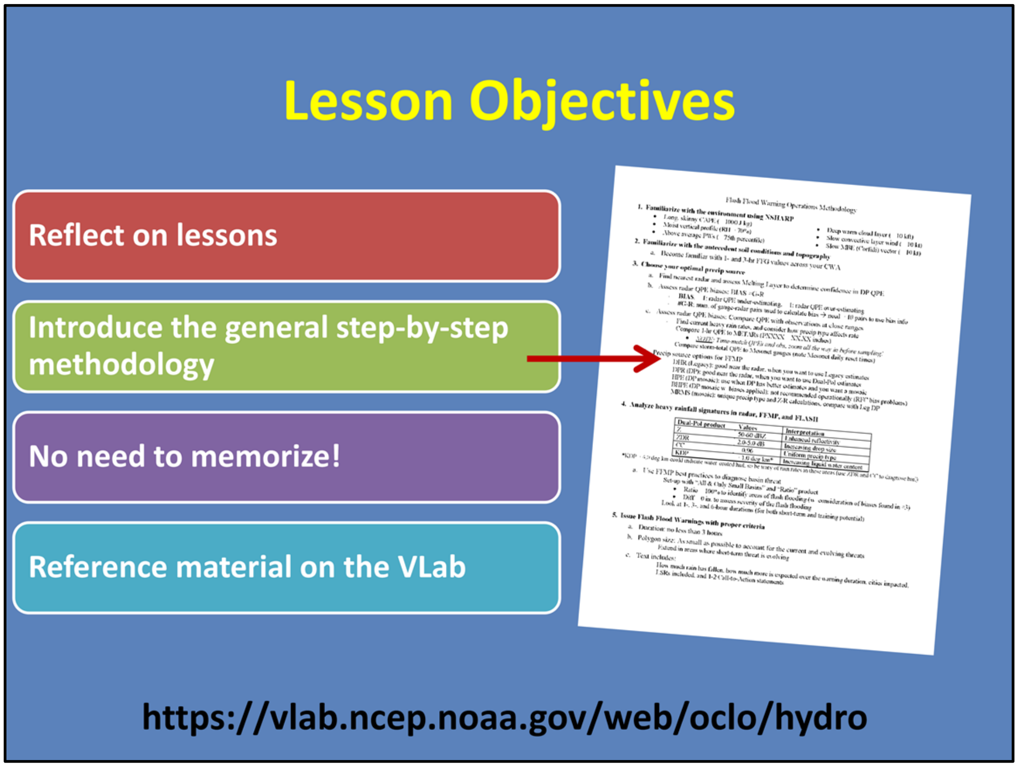

This module is different than most WDTD lessons because it's a chance for you to step back and reflect on the lessons that have led to this point. We'll tie them all together into one general step-by-step warning ops methodology, and we don't expect you to memorize this process. In fact, we have it all laid out for you on the VLab to reference at any time.

If you're taking this lesson as part of the Radar & Applications Course, you'll have the chance to apply this material soon enough in the Workshop Primer and workshop simulations.

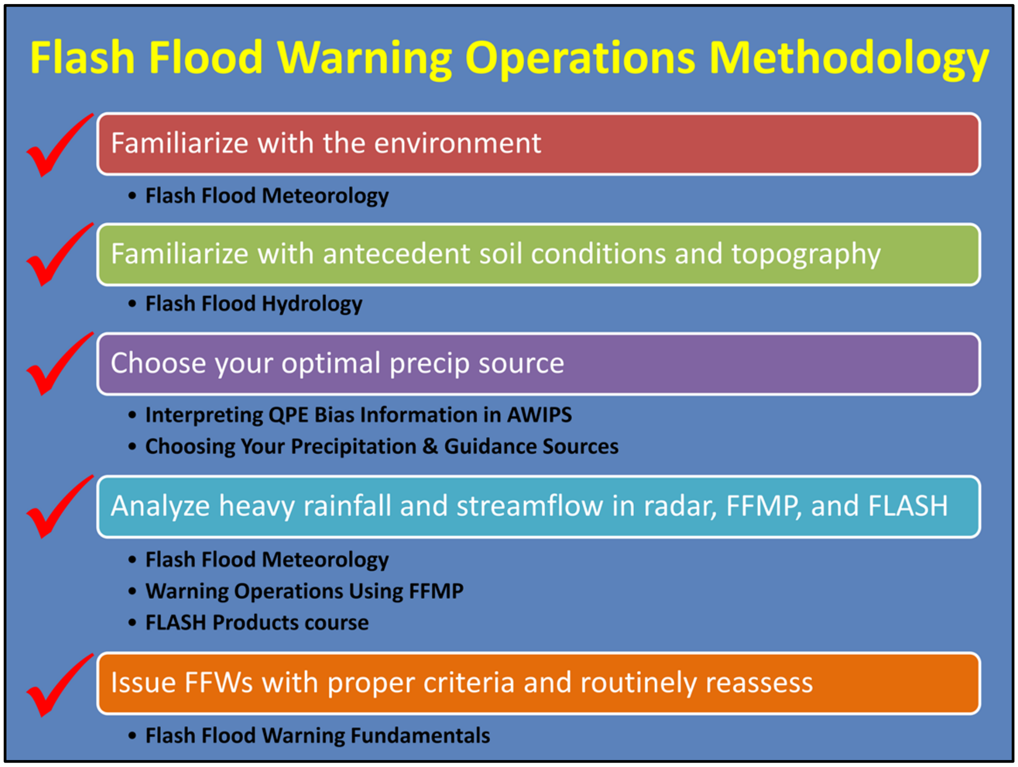

And here it is! This is the general process that we, at WDTD, think effectively aids in flash flood warning decision-making. While every office (and forecaster for that matter) will have differences when it comes to their hydro desk procedures, this step-by-step methodology is a good starting point.

It ensures you've: familiarized yourself with the current environment, antecedent soil conditions, and topography; are using the optimal precip source for the event; can analyze heavy rainfall and streamflow data via radar, FFMP, and FLASH; and are applying best practices when issuing warnings.

If you are taking this lesson as part of the RAC, congrats! You've already been introduced to each of these topics! But if you're taking this outside of RAC, please reference any of these WDTD lessons for more in-depth training.

The rest of this lesson will briefly summarize each of these steps.

Back to Top

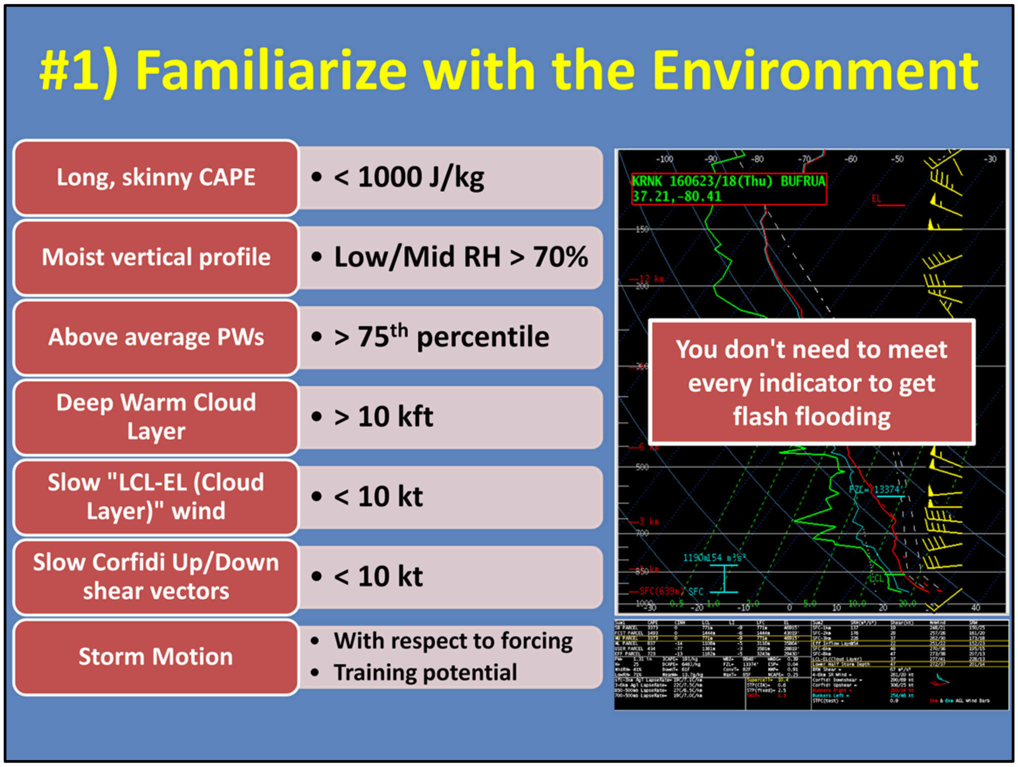

At the warning desk, one of the best ways to familiarize yourself with the environment is through an NSHARP sounding analysis. For flash flooding, some good indicators that the environment is primed for heavy rainfall are:
--a long, skinny CAPE profile (<1000 J/kg)
--a moist vertical profile (RH > 70%)
--above average Precipitable Water values (>75th percentile)
--a deep warm cloud layer (> 10 kft),
--slow cloud layer wind (< 10 kt), and
--slow Corfidi up/down shear vectors (< 15 kts)
--In your analysis, also consider storm motion with respect to a forcing mechanism and training potential.

Keep in mind that you don't need to met every one of these indicators in order to get flash flooding. You can have a high moisture day where any initiation is going to dump buckets, or you can have an average moisture day where the winds point to training storms being the main concern.

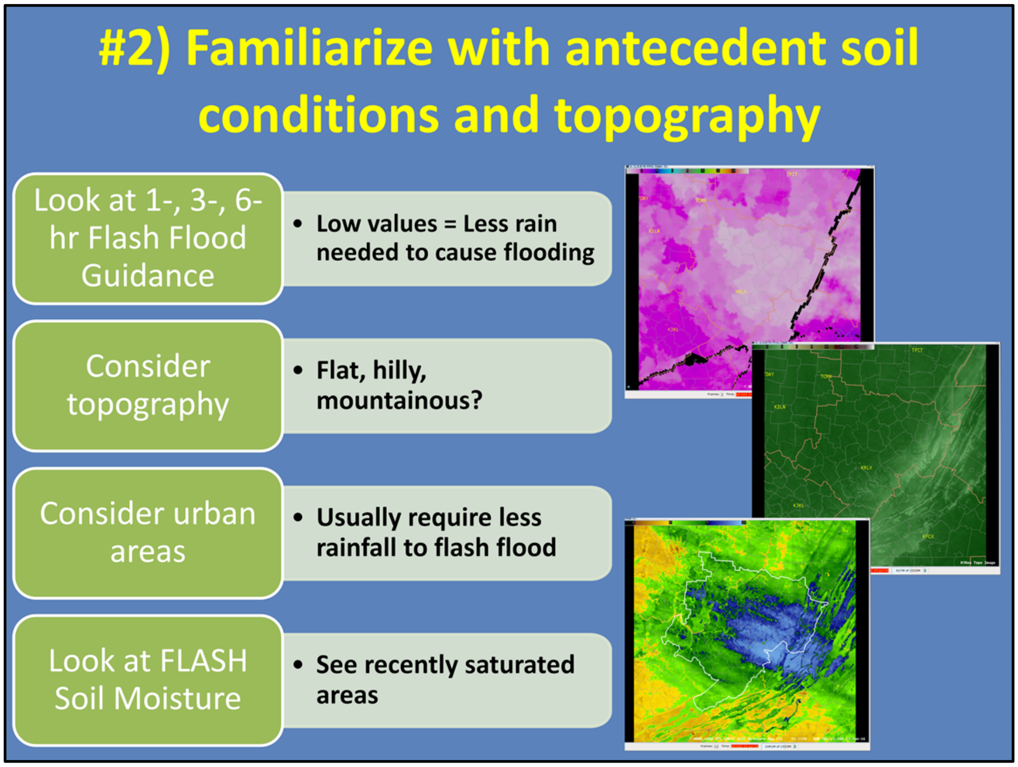

Next up is to familiarize yourself with the antecedent soil conditions and topography of your area. The easiest way to do this is using your Flash Flood Guidance products. For flash flooding, your 1-, 3-, and 6-hour FFG values will give you an idea of where recent rainfall may have already saturated soils. Remember, low values denote that less rainfall is needed for streams to overflow their banks. Keep in mind that these products are usually only updated up to 4 times a day, so if rainfall has occurred after the latest update, then it will not be reflected in the FFG products.

Topography also plays a role in where the water is routed during flash flooding, and how quickly it is routed. Also consider where your urban areas exist since they usually require even less rainfall to produce flash flooding.

Finally, if you use FLASH, consider each model's soil moisture product. This can help you see areas where FLASH has recently saturated soils and how that may affect model output.

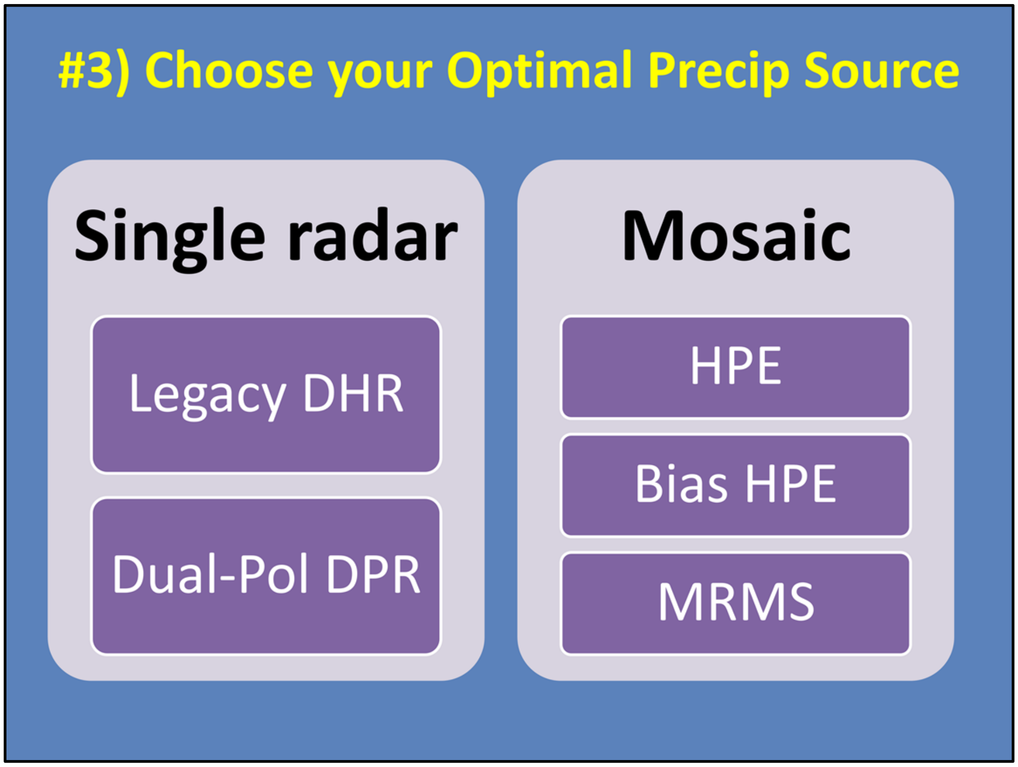

Step 3 is to choose the optimal precip source for use during warning operations, and it's not a trivial step. The best source can often change, so we have created some general guidelines to help you decide.

First, here is a list of what's available. You can learn more about the pros and cons of each one in the "Choosing Your Precipitation & Guidance Sources" lesson.

Back to Top

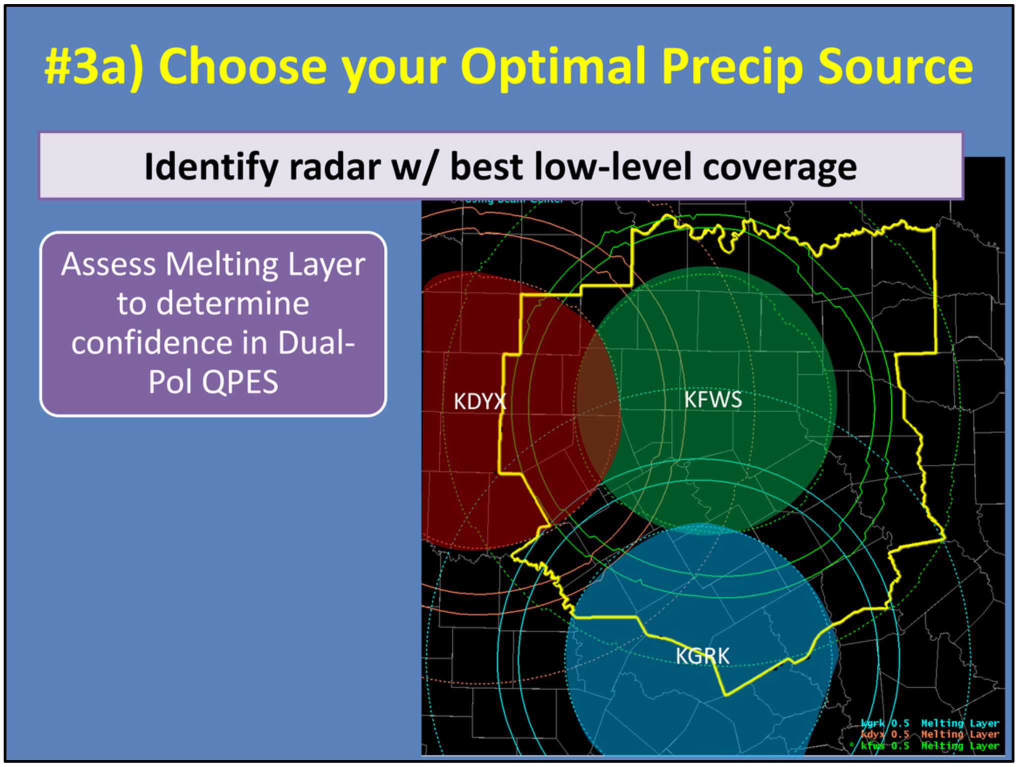

First and foremost, identify the radar with the best low-level coverage for the given storms. Keep in mind that this may not always be the *closest* radar, but usually that is the case. Here is an example of the Fort Worth CWA, with its 3 dedicated radars.

Assess the Melting Layer to determine where you can have higher confidence in your Dual-Pol QPEs. Your highest confidence is in areas that are below the Melting Layer, such as the green area of the KFWS radar. Within or above the Melting Layer, estimates could be affected by mixed or frozen precip classifications.

In this case, look at how much of the CWA isn't ideal for the KFWS radar. Depending on the location of the storms, using the surrounding blue and red radars may help you get the best QPEs.

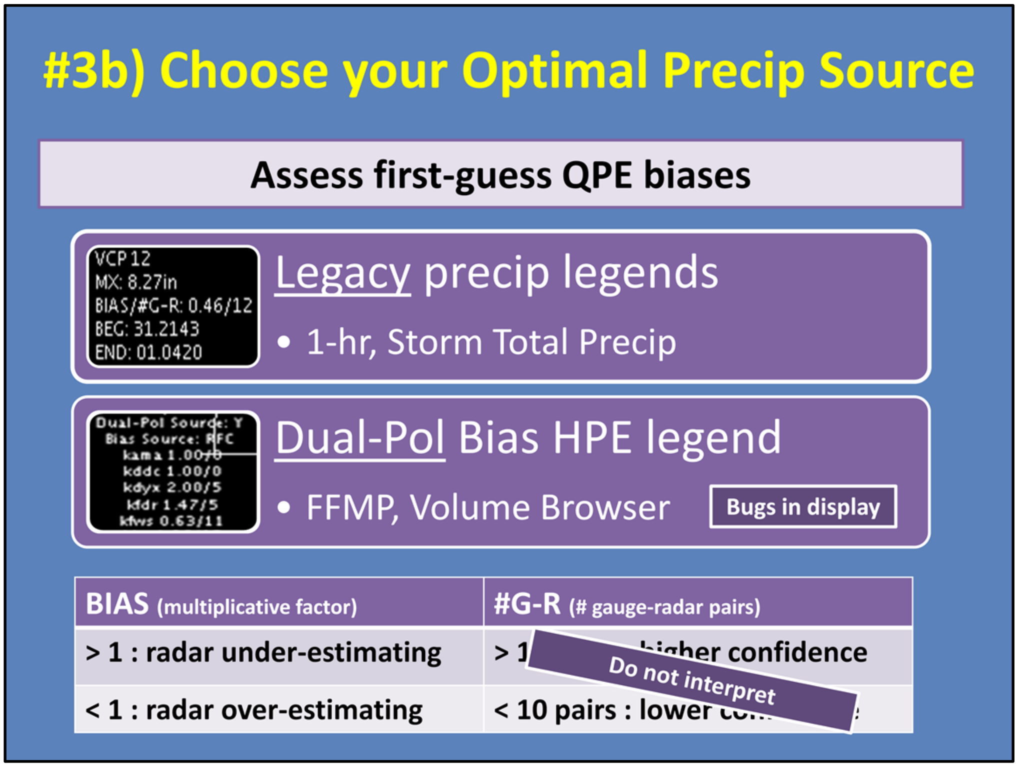

So once you have an idea of which radar is best, you can get a first guess of the potential bias of a precip source by reviewing its bias information. For Legacy, this is readily available via the 1-hour or Storm-Total Precip product legends, as seen here. For Dual-Pol, look at the Bias HPE legend in either FFMP or the Volume Browser.

As a reminder, the two values displayed are the bias factor itself, which can tell you if the radar is under- or over-estimating, and the number of gauge-radar pairs, which tells you how many pairs were used to calculate the bias. As of Fall 2018, do not interpret the gauge-radar pair information in these displays as there are bugs. Please review the "HPE and Bias HPE" lesson, as well as the "Interpreting QPE Bias Information in AWIPS" lesson for details and work-arounds.

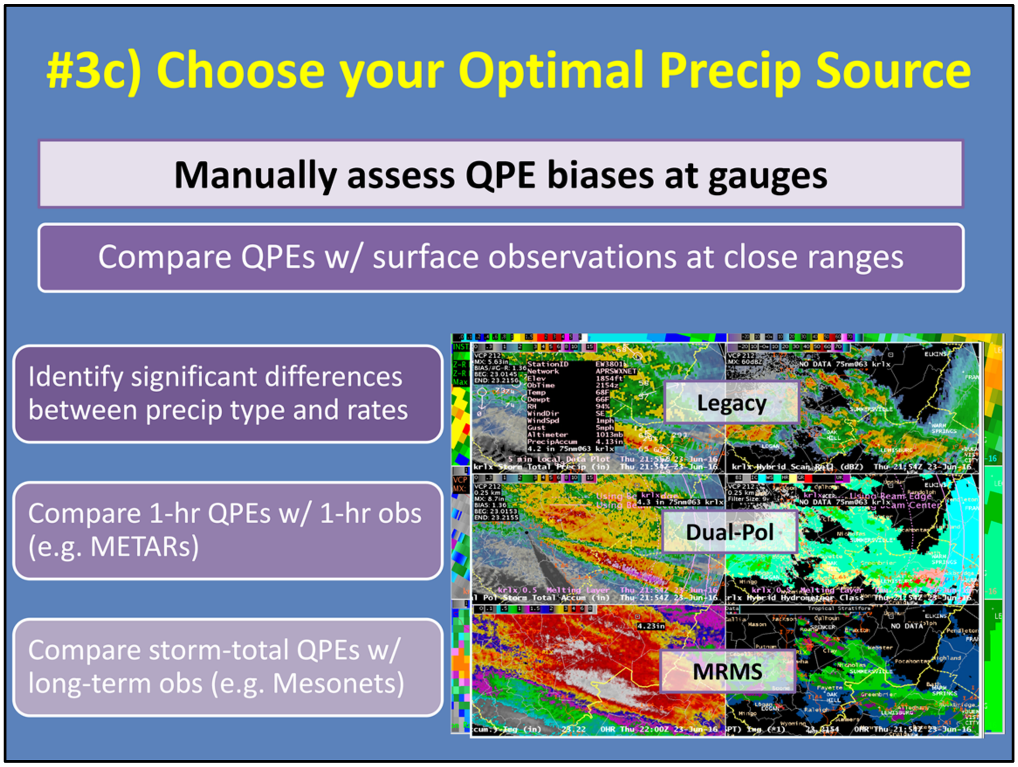

Probably the best way to get a feel for how each precip source is doing at any given time is to manually compare the QPEs with surface observations at close ranges. While gauges have been known to have their own issues, they are still the primary form of ground truth to calibrate yourself with potential precip source biases.

--Start by looking at Legacy, Dual-Pol and MRMS side-by-side to identify any significant differences between the precip types and their associated rates.
--Next, compare 1-hour QPEs with 1-hour observations, most likely through METARs. Keep in mind that you MUST remember to time-match in order to get a proper comparison.
--Finally, compare storm-total QPEs to longer term obs, like Mesonets, if you've got em. Get to know the local networks to know when these running totals reset, in order to make the best comparison possible.

Back to Top

#3) Choose your Optimal Precip Source

	Coverage	Dual-Pol?	Bias corrected?	Resolution	Accumulation products	Z-Rs
Legacy DHR	Single-radar	No	No (default, configurable at RPG)	1 km x 1 deg 3-6 min	1-, 3-, STP (and user-selectable)	Single Z-R, set at RPG
Dual-Pol DPR	Single-radar	Yes	No	0.25 km x 1 deg 3-6 min	1-, 3-, STA (and user-selectable)	Spatially-varying (based on HHC)
HPE mosaic	Mosaic	Yes	No	1 km x 1 km 5 min	1-hr	Spatially-varying (Inherited from DPR for each radar)
Bias HPE mosaic	Mosaic	Yes	Yes	1 km x 1 km 5 min	1-hr	Spatially-varying (Inherited from DPR for each radar, bias-corrected)
MRMS radar-only mosaic	Mosaic	No	No	1 km x 1 km 2 min	2-min update: 1-hr ≥ 1-hr update: all other accum	Spatially-varying (based on SPT)

Now put it all together to actually pick the precip source to use in warning decision-making. Is Legacy, Dual-Pol, or MRMS performing the best compared to obs? Will a mosaic help? What about bias corrections being applied? All of these factors should be considered to make your ultimate decision.

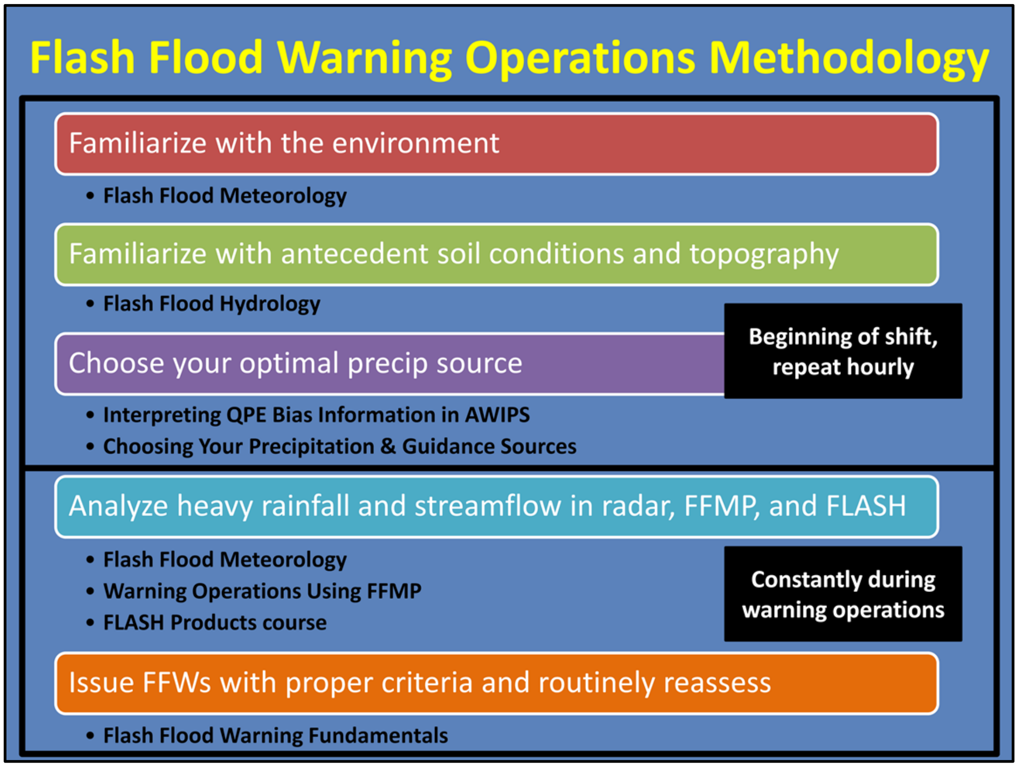

I'm going to step back for a moment and say...Steps 1-3 are ideally done at the beginning of a shift and/or beginning of an event. Once familiar with the environment, it shouldn't take much time to go back and repeat these steps each hour when new model runs, FFGs, and obs may be coming in.

As we move forward in our methodology, Steps 4 and 5 are then done continuously during warning ops. These steps are what you need to be able to do quickly and efficiently throughout your warning shift.

Back to Top

#4) Analyze Heavy Rainfall/Streamflow: Radar

Product	Values	Interpretation
Z	50-60 dBZ (40-55 dBZ tropical)	Enhanced reflectivity
ZDR	2.0-5.0 dB (0.5-3.0 dB tropical)	Bigger drop size (Smaller drop size)
CC	> 0.96	Uniform precip type
KDP	> 1.0 deg/km* (> 4.0 deg/km : water-coated hail?)	Increasing liquid water content

- Low-echo centroid signatures : precip below the freezing level
- Favorable supercell characteristics : slow, large updraft; moist inflow region

Once you have settled into a precip source and are ready to start picking out storms for warning decisions, then begin analyzing heavy rainfall and streamflow.

Here are the Dual-Pol characteristics that provide a lot of insight into where warm rain processes are dominating. Additionally, look for low-echo centroid signatures that show the majority of precip cores below the freezing level.

And don't forget that supercells can also produce heavy rainfall if they are slow movers and/or have the right environmental factors, such as a large updraft or very moist inflow region.

#4) Analyze Heavy Rainfall/Streamflow: FFMP

All & Only Small Basins

- **Configure "Layer" menu**

Start w/ Ratio, then to Diff

- **Ratio > 100% : location of the threat**
- **Diff > 0 in. : magnitude of the threat**

Duration

- **1-hr : short term**
- **3- and 6-hr : training storms**

Determine downstream direction

- **Configure "Click" menu**

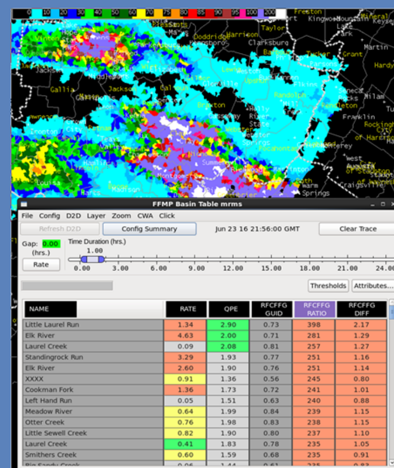

FFMP is a powerful tool that can help you slice and dice QPE and compare to FFG in order to diagnose flash flood threat. Here's some of the basics to effectively use FFMP:

--Always look at your smallest basins, since they are the most flash flood prone.
--Use Ratio AND Difference together to understand the location and magnitude of the threat.
--Consider 1-, 3-, and 6-hour durations for both short-term and training potential.
--And determine the downstream direction so you can anticipate where additional impacts could occur.

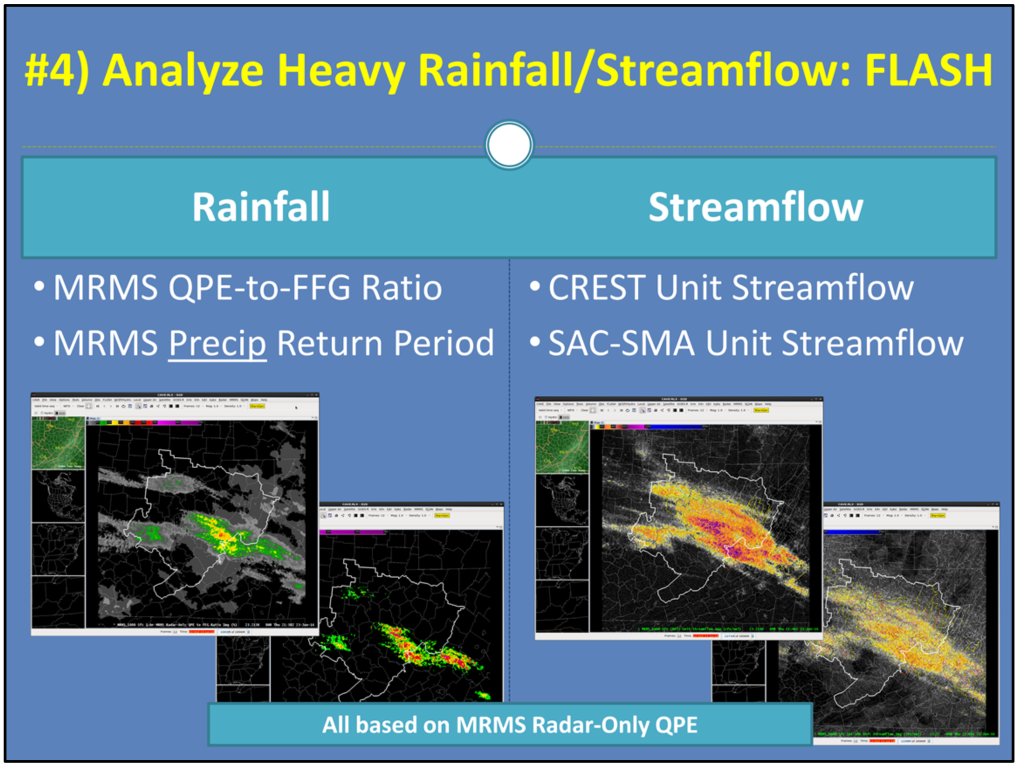

The FLASH suite of products is still new and its full applications are still being investigated. But there are some products that we know are useful now.

When you are interested in analyzing heavy rainfall, the MRMS QPE-to-FFG Ratio product gives a quick look at ratios similar to FFMP. And the MRMS Precip Return Period product compares the current rainfall to ARI thresholds. These gridded products are available at multiple durations.

When you want to analyze the hydrologic response, the unit streamflow products for both the CREST and SAC-SMA models are useful for diagnosing where above normal flow is occurring.

Keep in mind that all of these products are based around the MRMS Radar-Only estimates. So any biases in the QPE will immediately affect all of these products.

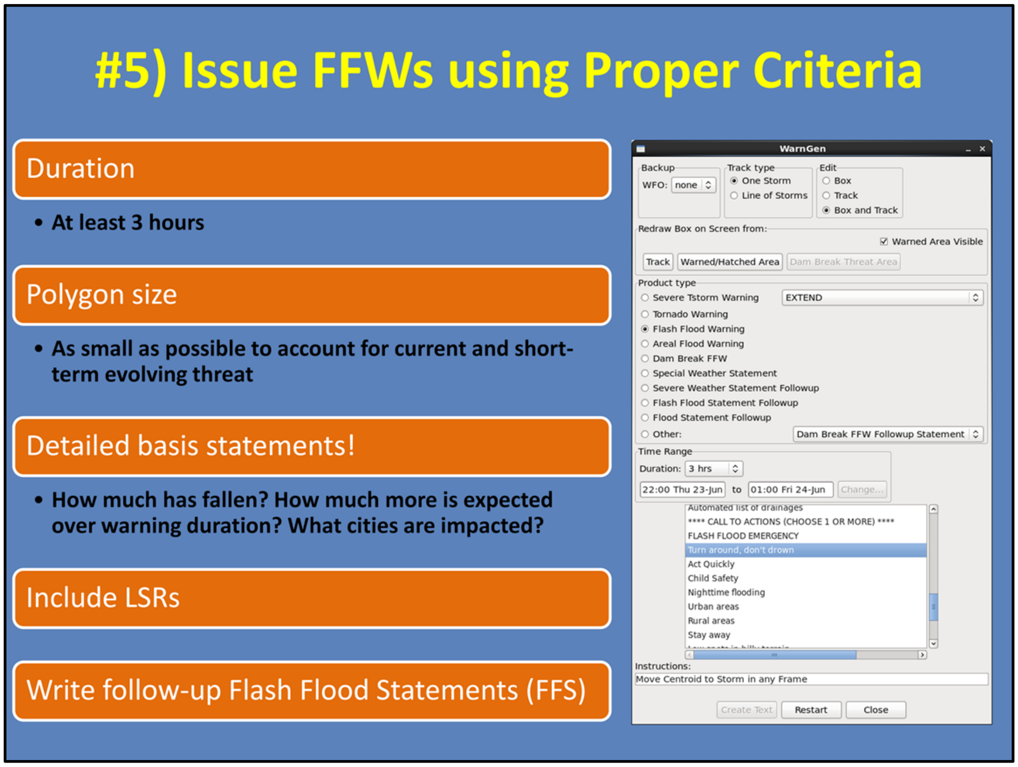

It all comes together in the final step of issuing a sound warning.

--Your warnings will generally be at least 3 hours in duration to account for rainfall, runoff, and receding time.
--Polygons should effectively cover the current and short-term evolving threat.
--The warning text should include relevant details about current and forecasted rainfall amounts and impacts.
--Always include Local Storm Reports if you've got em.
--Effectively communicate impacts through frequent updates, at least once per warning, when important reports arrive or information changes.

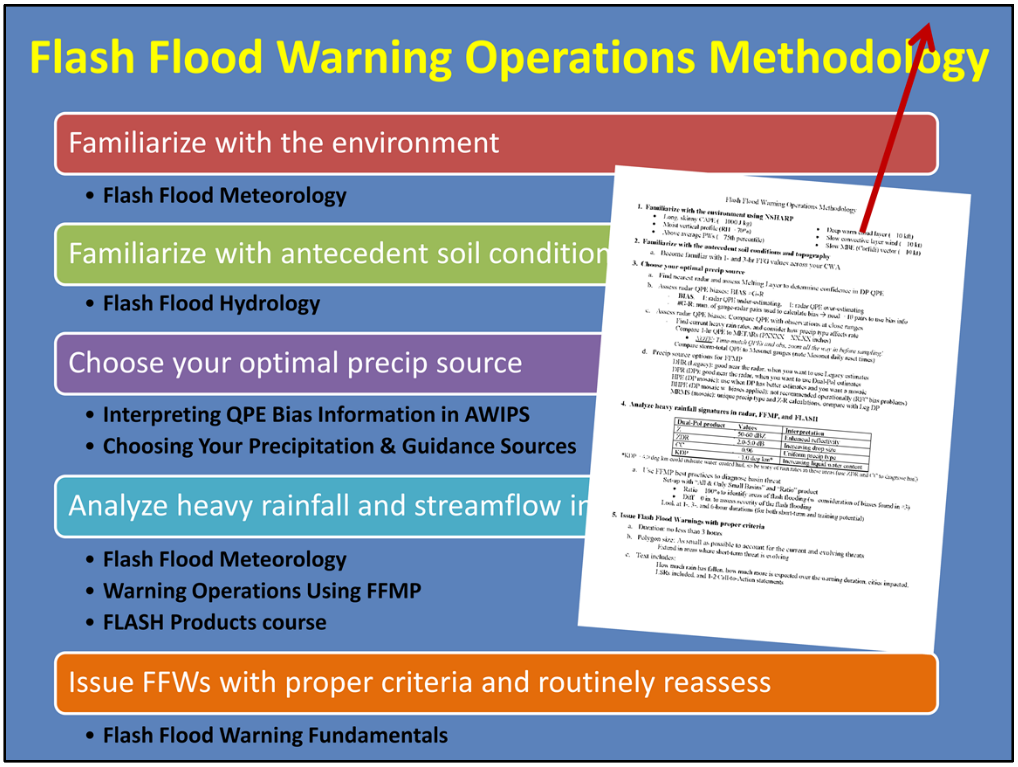

And there ya have it! Our Flash Flood Warning Operations Methodology in a nutshell. Remember to refer back to any of these lessons to get a better breakdown of each step. Additionally, we have created a one-page reference guide that summarizes all of this information in a printable form. It is available through the VLab page on the link in the Resources tab.

Thanks for taking this lesson! There is no quiz, so just close when you are ready.