Radar

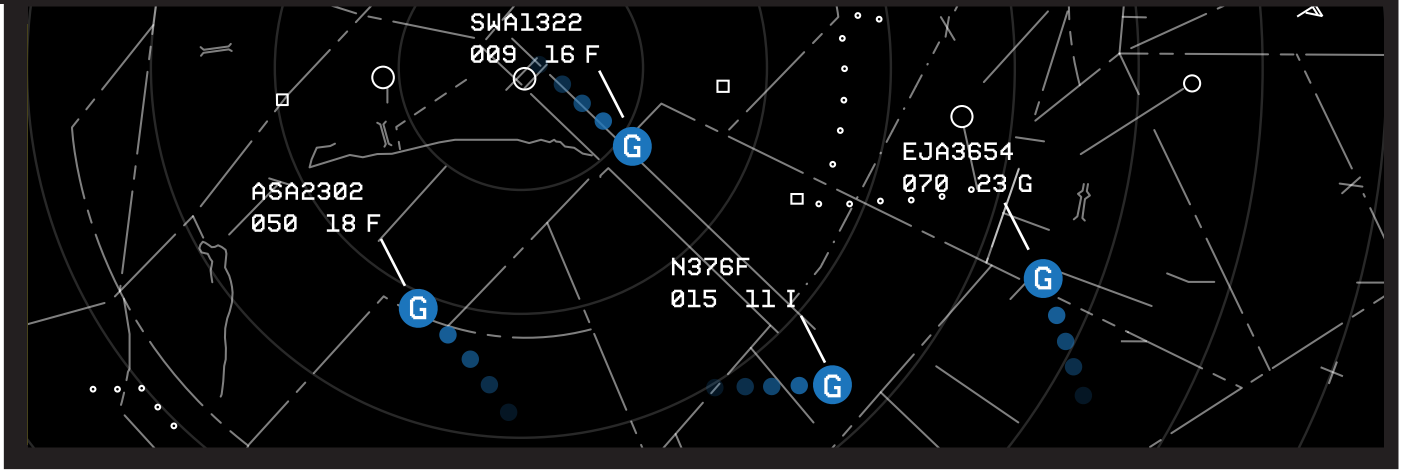

The radar scope is the heart and soul of air traffic control. It's a 2D representation that a controller builds into an ever-changing 3D picture of aircraft in the sky. Here's how that gets used.

How does ATC tell which target is me?

Identifying which target is you out of the dozens on my scope may seem like finding a needle in a haystack. However, imagine if a needle could tell you exactly where to look or make itself stand out amongst the other needles.

ATC uses two types of air traffic radar: *primary* and *secondary*. Our primary radar transmits electromagnetic energy as it rotates. This energy bounces off any aircraft (or other object) in range and returns to the radar antenna. It requires no interaction with the aircraft, but it shows us only your position and speed. We know *something* is out there, but not who or what.

Secondary Surveillance Radar (SSR) fills in the gaps. It's usually attached to the primary radar antenna so the antennas spin in sync. SSR interrogates aircraft transponders, which respond with their altitude, squawk code, and other information. It can determine an aircraft's distance by measuring the time between the interrogation and the reply, and relative position by the antenna's rotational position when the reply comes. However, it's blind to aircraft without a working transponder.

Our radar ID methods are divided into those two categories. Methods for Primary Radar Identification are laid out in the **7110.65 5-3-2**.

"Observing a departing aircraft target within 1 mile of the takeoff runway end at airports with an operating control tower …" As I'm working Tower, you're a flight requesting radar services, and I'm about ready to clear you for takeoff. I need to notify the radar controller that your aircraft will be my next departure. Local procedures vary on how this is done, such as verbally, by making your call sign flash on their scope, or by sending them a strip electroni-

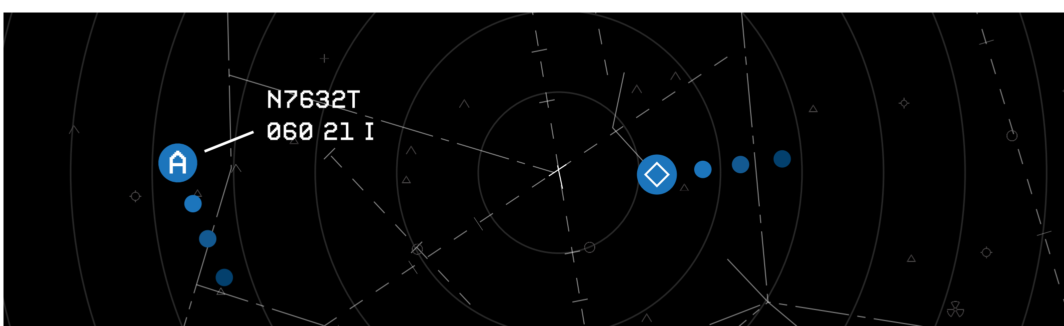

Here's how things might look on an Approach controller's radar scope. The hollow diamond is a primary target (no transponder signal). The other shows a datablock that can contain aircraft ID, altitude, airspeed, etc. Some of this comes from Mode-C, some can come from ADS-B Out, and some can be added by the controller at the scope.

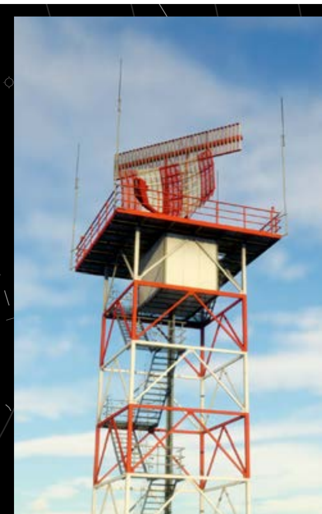

Pilots have their own version of "radar" now with ADS-B targets showing on their display of choice in the aircraft. Some of the data is the same, but some is reversed. The pilot sees relative altitudes above and below and a trend vector showing where the aircraft is going. ATC sees the pressure altitude automatically corrected to a local altimeter setting and a trail of where the aircraft has been.

cally or physically (via a drop tube). When your target pops up off the end of the runway, they're already expecting you and can identify you with confidence.

"Observing a target whose position with respect to a fix … has been accurately determined and made available to the controller …" If you tell me your location relative to an aviation fix, geographical point, or airport, I can find you. Keep it simple: just your mileage and bearing from the fix. Examples? **"I'm 10 miles north of Crater Lake at 9500,"** or **"I'm over Sioux Falls Airport at 6500."** If you're the only target at that location and altitude, I can positively identify you.

"Observing a target make an identifying turn or turns of 30 degrees or more …." If there are multiple aircraft near your stated position or there's a question about your position, I'll have you make a continuous left or right turn, and watch which target makes the same turn. This could take a little bit of time, and even include another turn in the other direction to be certain.

Secondary radar identification techniques require you to interact with your transponder when requested. They're covered in the **7110.65 5-3-3**: Beacon/ADS-B Identification Methods.

"Request the pilot to activate the 'IDENT' feature of the transponder/ADS-B and then observe the identification display." You push IDENT and your target flashes on an Approach scope or shows as three horizontal lines (a "triple slash") on a Center Scope. Either way, it draws my attention to it. Controllers often pair it with the next method.

"Request the pilot to change to a specific discrete or non-discrete code, as appropriate, and then observe the target or code display change." When you request either VFR flight following or an IFR clearance, I'm going to assign you a squawk code. Upon seeing your target's code change to the new code, I can consider you identified.

"Request the pilot to change their transponder/ADS-B to 'standby.' After you observe the target disappear for sufficient scans to assure that loss of target resulted from placing the transponder/ADS-B in 'standby' position, request the pilot to return the transponder to normal operation and then observe the reappearance of the target." Similar to the previous tactics, I'm looking for a change in your target status. As you switch your transponder to standby, your target will revert to a primary-only target. Once you turn it back on, your altitude will be visible again. Observing these instructed changes allows me to identify you.

In reality, I'm often using a combination of these methods. Imagine you make the following request: **"Approach, N123AB is a Cessna 172, 10 miles east of Harrison VOR, at 5500 feet, requesting flight following to Fayetteville."** I type you into the system, which generates a squawk for you, and then say, **"N123AB, Approach, squawk 4652, IDENT."**

Meanwhile, I've already been looking at the only target around 10 miles east of HRO VOR, indicating 5500 feet. Suddenly, that target starts flashing, changes to your assigned squawk, and transforms into a full data block with your call sign. Consider yourself triple-extra-identified: position correlation, IDENT, and a squawk code change.

What else am I doing when I identify you? I'm making sure that you're actually in my airspace. If not, I'll get you talking to the correct controller.

What is IFR radar separation?

Before we can talk about radar separation, we need to answer this question: What kind of radar are we talking about?

Air traffic controllers have several sensors available to them. Traditionally, tower and approach control facilities relied on short-range Airport Surveillance Radars (ASRs) that could see out to about 60 miles. Centers used Long-Range Radars (LRRs) with an effective range of 200-250 miles. These radar sites (both of which contain primary and secondary radars) have been supplemented by hundreds of ADS-B ground stations nationwide.

Our IFR radar separation requirements depend on which of these sensor(s) are being used to track your aircraft. Many moons ago, at my first approach control, we had a single ASR. No other sensors. If aircraft were within 40 miles of the antenna, my minimum separation between IFR aircraft was 3 miles. I needed 5 miles for aircraft farther out.

Things have changed. Tower and Approach now use a system called FUSION that seamlessly blends ASR, LRR, and ADS-B sensors. This overlapping coverage allows terminal controllers to continually use 3 miles of IFR separation in most cases.

At centers, their historical reliance on slower-spinning (and therefore slower-updating) LRRs traditionally required 5 miles of IFR separation. That's still the case for aircraft at altitudes above FL230. However, since the adoption of the ERAM (EnRoute Automation Modernization) computer system in Centers, access to combined ASR, LRR, and ADS-B feeds allow them to use 3 miles of separation in certain areas defined by local procedures. These areas must be within 40-60 miles of a preferred radar site.

If a certain sensor feed fails, or goes offline for maintenance, the coverage lapse may require increased spacing. FUSION would advise me of this by adding "ISR" to an affected target's datablock, meaning "Increased Spacing Required." A note in **7110.65 5-5-4** says, *"In the event of an unexpected ISR on one or more aircraft, the [controller] working that aircraft must transition from 3-mile to 5-mile separation, or establish some other form of approved separation as soon as feasible."*

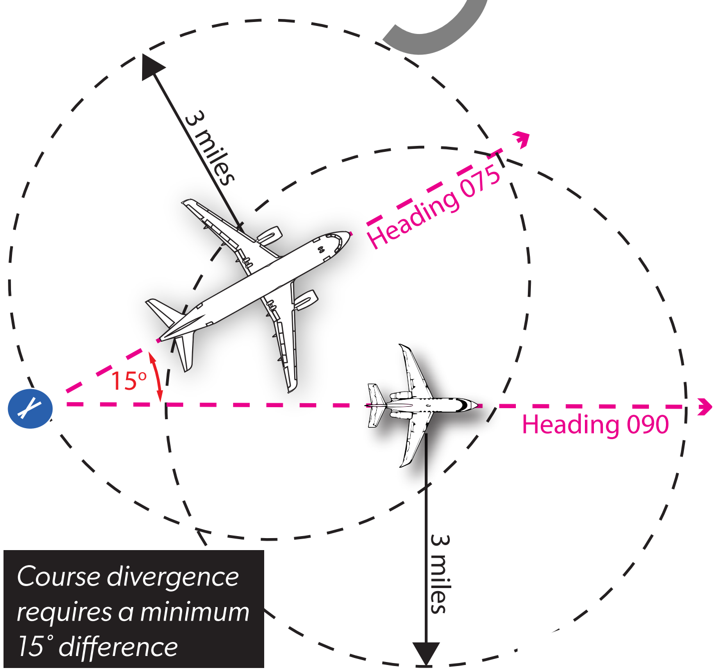

Course divergence requires a minimum 15˚ difference

Of course, we're keeping aircraft away from obstructions as well as each other. If you're IFR, flying near an antenna, and all my FUSION feeds are working properly, I must keep you a minimum of 3 miles from an obstruction. If one of my radar feeds goes into the red and "ISR" starts blinking on your target, that becomes 5 miles.

I also need to separate you from other radar sectors or airspace. In the approach and tower environment, I can't let an IFR get closer than 1 ½ miles to another controller's boundary without coordinating with them. Centers must maintain a minimum of 2 ½ miles of boundary separation. These standards can increase depending on your distance from the antenna.

Those numbers are just the baselines. Get deep into the weeds of the **7110.65's** radar separation regulations and you'll find things like flights of multiple aircraft getting an additional mile of separation, or that we must apply wake turbulence distance minima, when necessary.

How much separation do IFR and VFR aircraft need?

You're IFR in a cloud layer at 3000 feet when I call traffic to you: **"N123AB, traffic, 1 o'clock, 3 miles, westbound, 2500, Piper Cherokee, converging course."** You hear me call you out as traffic to the Cherokee as well.

It's a total whiteout beyond your windshield. **"Approach,"** you respond, **"N123AB is IMC. Negative contact."** You've told me you can't see the VFR traffic and—with you in the clouds—it's highly unlikely the other pilot will spot you. No "see and avoid" today. How much space must I ensure between you and the VFR traffic?

*There's another type of IFR radar separation: **course divergence.** Approach controls and towers can allow two aircraft to get closer than 3 miles when (1) the course of one aircraft is projected to pass behind the other, (2) their targets will not touch, and (3) the aircraft's courses diverge by at least 15°.*

It depends on the type of airspace you're currently occupying.

Let's start with Class C. Per **7110.65 7-8-3**, I need either visual, vertical, or lateral separation. My vertical separation requirement is what I've already got: 500 feet between the VFR and you. I just need to maintain it. I tell the VFR to remain at or below 2500 feet and everyone's clean.

What if you were both at or near the same altitude? In the lateral plane, I only need target resolution between a VFR and an IFR aircraft. Translation: *I can't allow your radar targets to touch.* If you're thinking, "Whoa, if that VFR's target gets anywhere near close to rubbing against my target, that sounds bad," you are correct.

While a VFR target can legally swoop within a few pixels of your IFR aircraft on my scope, it would be irresponsible to let you get that snuggly. I prefer to build an early margin of safety by using vectors or separating you by altitude. The **.65's 5-5-3** also says, *"Mandatory traffic advisories and safety alerts must be issued when [Target Resolution] is used."*

In Class B airspace, the vertical and lateral standards increase slightly if the VFR aircraft is operating near larger or jet aircraft, regardless if that aircraft is VFR or IFR. Per **7110.65 7-9-4 (b)**, *"VFR aircraft must be separated from VFR/IFR aircraft/helicopter/rotorcraft that weigh more than 19,000 pounds and turbojets ..."* by no less than 1 ½ miles separation, 500 feet vertical separation, or visual separation.

Even in Class C or B, if the aircraft can maintain visual separation from each other, then that's a different animal. The vertical and lateral separation doesn't apply any longer because see-and-avoid rules are in effect.

What about aircraft flying outside the confines of Class C or B? There aren't any strict separation requirements between VFR and IFR aircraft (and certainly not between two VFR aircraft). There are a ton of unidentified VFRs flying around at any given second, not talking to ATC, that should be "seeing and avoiding" other traffic. Theoretically. With their type, altitude, and intentions unverified, I also can't guarantee 500 feet of vertical separation, or 1 ½ miles, or any of that from aircraft not on my frequency.

That doesn't mean I just throw up my hands and bet on a solo student pilot to see you while in the midst of a power-on stall. My stated primary purpose is preventing a collision, and that doesn't distinguish between IFR or VFR aircraft. I must accomplish it using whatever method is most appropriate at the time.

Why doesn't ATC always give me the published missed approach?

In a word: Expedience. Published missed approach procedures are focused primarily on separation from surrounding terrain and obstacles, but they can have a significant impact on other air traffic. Also, when you're requesting multiple approaches, alternate missed approach instructions may be the best option to get you set up for your next request.

Let's first talk about traffic flow. Imagine I'm working Approach, overseeing a sea-level airport with both ILS and RNAV approaches to Runway 9. There's a VOR on the airport. The published missed for the ILS is an initial climb to 1000 feet, a climbing right turn to 2500 feet to intercept the VOR's 180° radial, culminating at a holding fix due south of the airport.

The published missed for the RNAV, in contrast, features a straight-out climb to 2500 feet, then a right turn to follow a series of waypoints ending at the same holding fix used by the ILS. In overview, they are similar: climbs to 2500, right-hand turns, same holding fix.

When there's known weather near an airport and an increased chance of a go-around due to wind shear or visibility, ATC may assign you (or you may request) alternate missed approach instructions to steer you toward a safer direction. I'll typically phrase it: **"In the event of a go-around, fly [heading] and maintain [altitude]."**

However, while I give approval to most aircraft requesting the published missed for the ILS, I'll regularly deny it for the RNAV approach. Why would I do that?

Usually, the aircraft requesting these approaches are single-engine piston trainers, like Cessna 172s or Piper Cherokees. A realistic average climb rate for a C172 is about 500 FPM. When a C172 goes missed on the ILS, it'll take at most a couple of minutes for it to hit 1000 feet and start its turn out of the way. Not too terrible.

Today, though, you're in a Cessna 172 requesting the RNAV approach with the published missed. If I approve it, you'll need to climb all the way to 2500 feet before turning. That's realistically a full 5 minutes that you'll be in my departure corridor, clawing for altitude at 80-ish knots. Meanwhile, I've got a line of jets I normally launch every 60 seconds just waiting to go. Every minute your Skyhawk is in the departure path is one departure that didn't get out.

To prevent this, I'll assign you alternate missed instructions. **"N132AB, at the completion of your approach, fly heading 180. Maintain 3000. Departure frequency 118.5."** Different controllers have different ways of phrasing this. For example, some say **"missed approach instructions, fly heading …"** or **"on the go, fly heading …"** I've taken to stressing the **"at the completion of your approach,"** since I've had some pilots interpret climbout instructions as a heading and an altitude to fly *right now.*

Alternate missed instructions at towered airports are usually standardized in SOPs or LOAs, with specific alternate missed instructions for each runway. That's what I'll issue you. Since both the tower and approach control are aware of these procedures, they require little to no coordination.

Once you complete the low approach and start climbing, I'll be watching for you to start your turn at the appropriate time. If you are late doing so, I'll remind you with, **"N123AB, execute your turn."** As you start turning, I start launching my jets. While I denied you the full missed approach for traffic flow, I haven't forgotten your initial request. I can still offer you the published hold. (The hold was never the issue, just the turn to get there.)

Weather also impacts the approval of published missed approach procedure. Let's say this time the ceiling is OVC005 and you want the ILS Runway 9 with the published missed. So you request that—and I come back with **"Unable."** Why?

The published missed has you hitting 1000 feet, followed by a climbing right turn to the VOR's 180° radial. Does it specify a heading to intercept that radial? No. You could hook a right all the way to a 270° downwind heading to join the radial.

That's bad because you'll be head-on with any IFR traffic I have on final. I can't use tower-applied visual separation between you and the other traffic when you'll be swallowed by the overcast at 500 feet.

To ensure separation, I assign you alternate missed approach instructions to diverge you away from my final. As long as your course diverges from the other aircraft by between 15 and 135 degrees, you're legal. *(See "Course divergence requires a minimum 15° difference" on page 61.)*

My final traffic to Runway 9 is flying a 090 heading. Take that 90°, add 135°, and you'd get 225°, which is the furthest right I can let you turn and maintain IFR separation. Do you know what already complies with my divergence requirement—and doesn't require any special coordination? Assigning you the standard alternate missed: 180 heading, 3000 feet, and 118.5 for the frequency.

Alternate missed instructions can set you up for the next thing. After this ILS Runway 9, you want to fly the RNAV to Runway 18, with a circle-to-land on Runway 9. Could I give you my standard 180° heading? Sure, but that's the opposite of where you need to go. Instead, I'll assign you a northerly heading. The catch is that I must coordinate it with the tower, since it's not one of our standardized climbouts. I make the landline call: **"Tower, Approach, approval request, N123AB on the go heading 030, up to 3000."** The tower checks the traffic for any potential conflicts and responds, **"Approved."**

Now I can tell you, **"N123AB at the completion of your approach, fly heading 030. Maintain 3000. Departure frequency 118.5."** It saves us both from expending more time, vectors, and gas than necessary.

TIP If you must fly a published missed approach for training, try to be flexible about it. Try to pick a time when the airport of choice isn't busy.

Anatomy of a STARS Radar Display

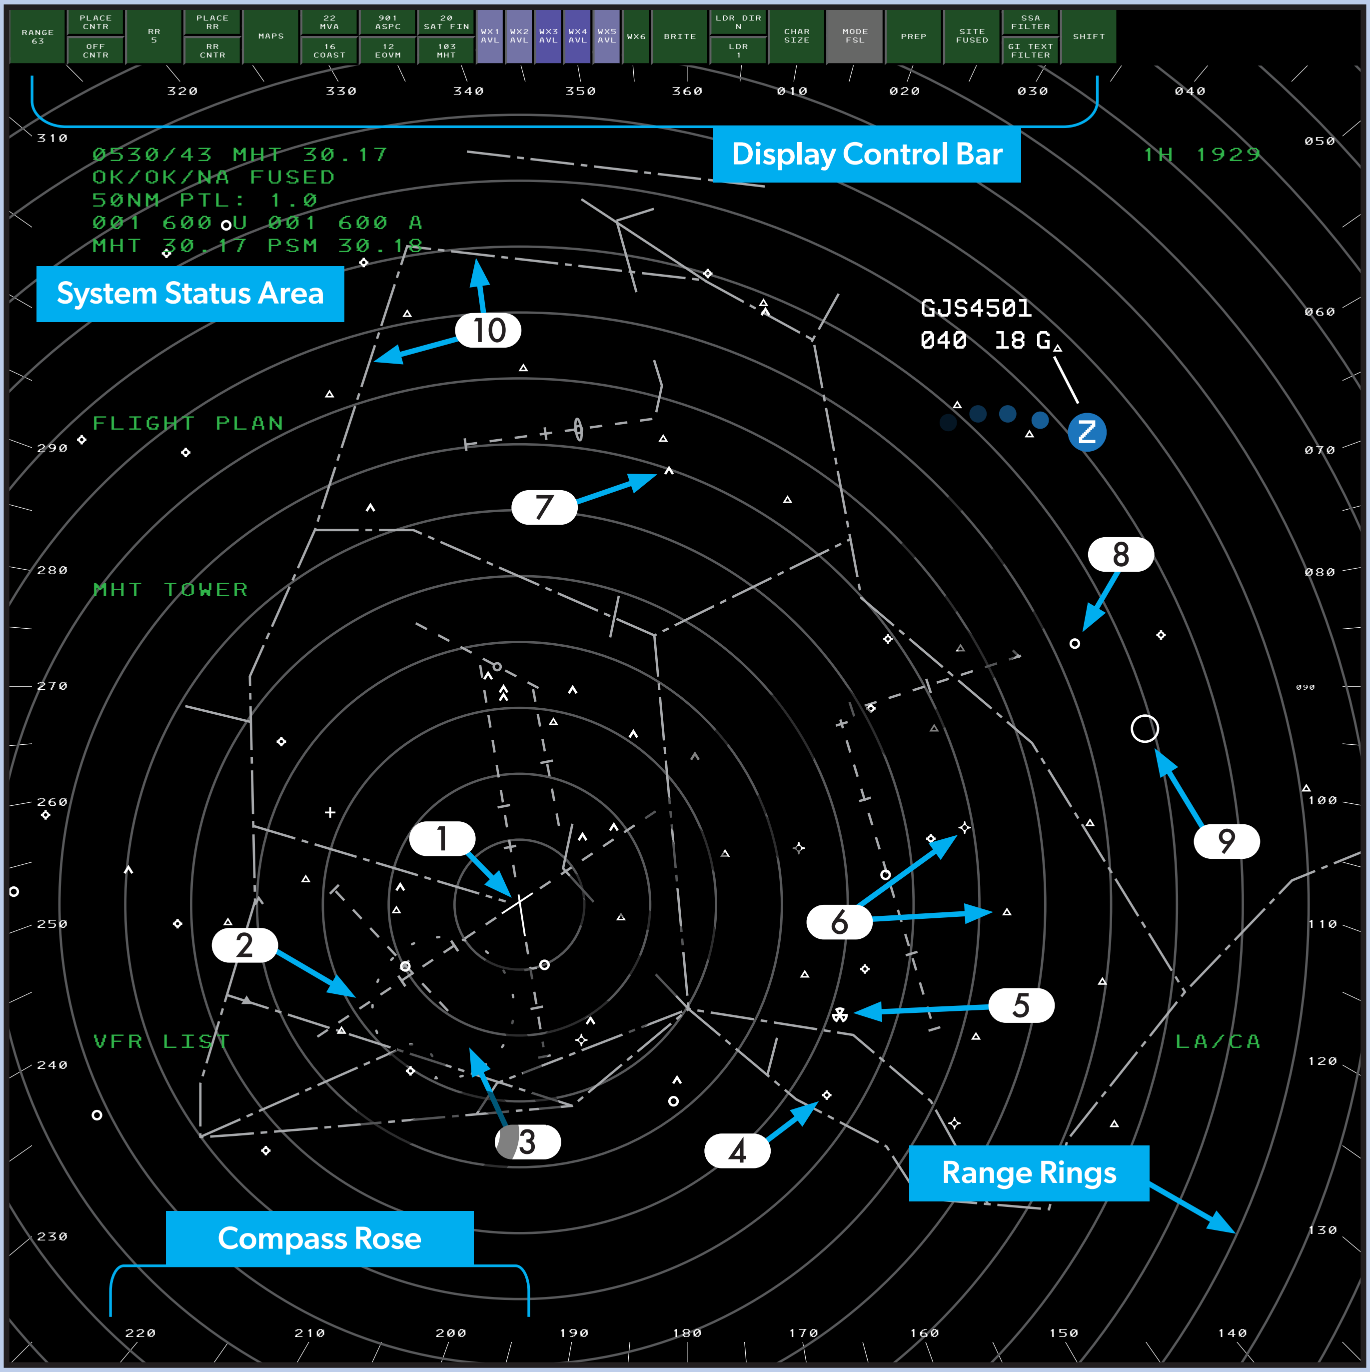

A radar display may look like a random collection of lines, shapes, and text at first glance, but everything on it has a specific purpose. Every element is designed to convey as much information as possible with minimum complexity and distraction. With training, a seasoned controller can look at a scope for just a few seconds, and get a grasp of what every target is doing and what kind of weather conditions are impacting their traffic.

Scopes are also quite customizable. Controllers each have their own techniques and working style. We're able to build and save our own preference sets for font sizes, map brightness, map selection, weather intensity levels, range rings, and more.

When I relieve another controller, the first thing I do is punch in a couple quick keystrokes and—boom–the scope is instantly configured just the way I like it. I have a different PrefSet for each sector or tower configuration I work.

Anatomy of a STARS Radar Display

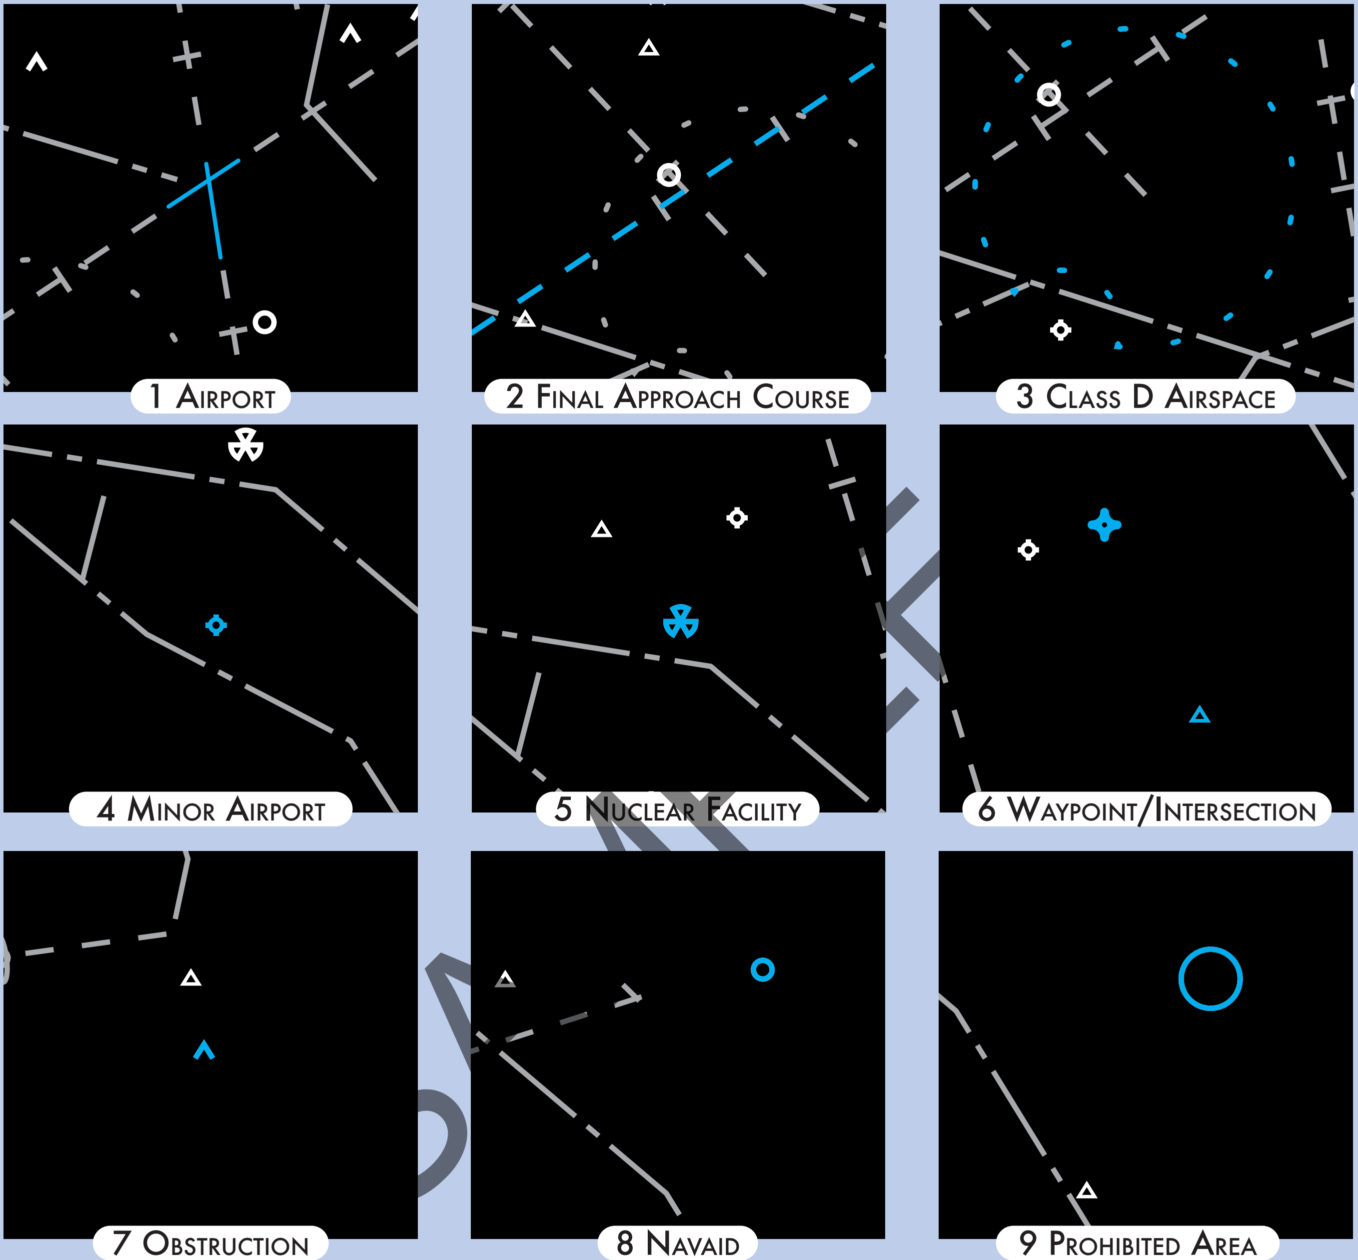

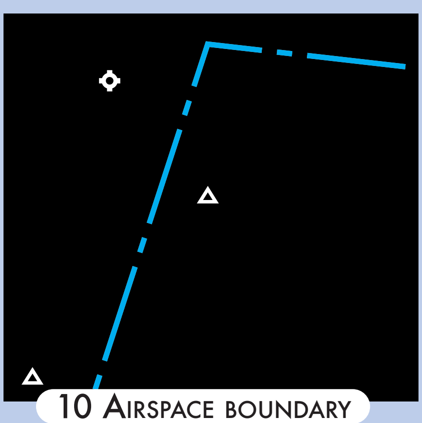

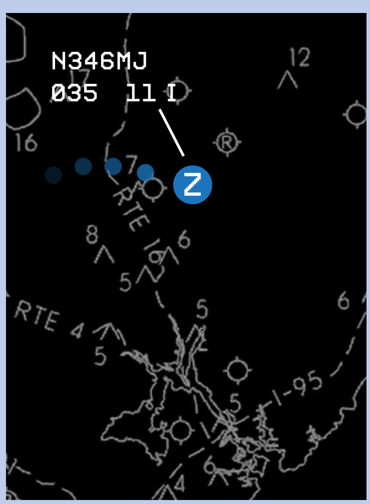

I also have specialized maps available. If I have an emergency or I need geographical reference, I can select my emergency map, which displays major highways and roads, as well as the shoreline of major bodies of water, rivers, and additional topography.

Some facilities have maps that indicate county borders (useful for contacting law enforcement regarding a laser or drone sighting). Facilities that work MOAs and other SUA may have maps customized to their needs.

ATC didn't give me alternate missed instructions for my practice approach. Do I fly the published missed?

I try to ensure every practice approach aircraft has alternate or published missed approach instructions. Controllers could be running a multitude of similar practice approach aircraft, into similar airports, requesting similar things, and each with similar markings on the scope. We each have techniques for tracking who's been told what.

Nevertheless, I'm not perfect. The law of averages means an oversight could occur. For instance, if I've cleared the same plane for two ILS approaches with the same climbout instructions at the same airport and am vectoring it for a third ILS (that will also get the same climbout), I may mistakenly believe I already issued them a climbout for this latest approach.

This is where you can step in. If you're expecting alternate missed approach instructions and haven't received them, please request them. After issuing an approach clearance, I've told a pilot, **"Contact Tower,"** or **"Change to advisory frequency approved,"** only to get the response, **"Do you have climbout instructions for us?"** Alternately, I've switched an aircraft to Tower, only to receive a landline call from the tower controller. **"Approach, Tower. N123AB says they need climbout instructions. What would you like for them?"** The pilot clearly asked the tower to confirm climbout instructions.

That verification works great if you're in active contact with a radar or tower controller. When approaching a non-towered airport, there's no tower controller to ask and (depending on radio coverage) you may be unable to communicate with me. What now?

The resolution depends on whether you're IFR or VFR. If you're IFR, then **14 CFR § 91.123** applies. *"Except in an emergency, no person may operate an aircraft contrary to an ATC instruction in an area in which air traffic control is exercised."* If you have been cleared for an instrument approach under IFR (an ATC instruction) the published missed approach procedure is 100 percent part of that approach. Without alternate ATC instructions, that's what you should execute. Remain on your squawk, fly the published missed, and contact ATC on the published frequency.

If you're VFR, then **AIM 4-3-22 (e)** is applicable. *"VFR aircraft practicing instrument approaches are not automatically authorized to execute the missed approach procedure."* A missed approach procedure must be specifically requested by you and approved by ATC. That didn't happen here. *"Where no separation services are provided during the practice approach, no separation services will be provided during the missed approach."*

If you were told, **"Practice approach approved. No separation services provided. Maintain VFR,"** and don't get climbout instructions,

Separation During VFR Practice

If you're flying under IFR, even in VMC, you get IFR separation throughout the entire procedure including the missed approach.

Under VFR (in VMC of course), there are two options. If there's a written procedure outlining how VFR practice approaches are handled (typically at towered airports), you will usually get separation services once the approach clearance is issued. It's only 500 feet of separation, but it's still provided.

That's partly why you must request to fly the published missed approach, if you want it, and ATC must approve. You're VFR and could, in theory, fly wherever you want, but you're working with a controller who is providing a service separating you from some other aircraft. We have to work together on this one.

You could also be flying your approach at an airport without existing ATC procedures for practice approaches. In this case, you typically won't be vectored onto the approach and will only hear, **"Practice approach approved. No separation services provided. Maintain VFR."**

That's typically at a non-towered airport. Very few controllers would be willing to provide IFR separation to a VFR aircraft at a non-towered airport. We have no idea what's going on in the traffic pattern.

This is also why requesting alternate missed approach instructions at a non-towered airport to efficiently move your flight along to your next practice airport may not be possible. The controller might be fine with something in a flat area. If there's terrain around, the controller might not be able do that without something that's written down.

Nearness doesn't mean, uh, "firstness." ATC needs to consider each aircraft's flying miles, speed, and other factors to build a sequence. Your close proximity to an airport doesn't guarantee you'll beat other traffic to the runway.

then there are no specific restrictions on how you can maneuver after the approach. Just maintain VFR, look out for other traffic, and contact ATC after the approach if you want additional services.

I was near the airport. Why did ATC make me last in their sequence?

Imagine you're IFR in a Cessna 172. I'm vectoring you for a visual approach to Runway 9. I've got you on a 5-mile midfield right downwind. The day's so clear you can just about read the Daily Specials board at the FBO restaurant. **"Approach, N123AB, I've got the field in sight,"** you report.

"Roger, N123AB," I reply, **"you're number three, following jet traffic from the west."**

Number three? But the airport's *right there.* What gives?

When mixing and matching aircraft of different types in a sequence, I take into account multiple factors. I look at speed. I ensure I'm complying with any LOAs with the tower and with IFR separation guidelines. And I look at their *flying miles*—not necessarily distance—from the airport.

What are flying miles? Literally, they're how many miles you need to fly before you get where you're going.

Right now you're westbound, 5 miles due south of your destination airport. However, I'm vectoring you for a (pretty standard) 5-mile final. You would have to fly 5 miles west on the downwind, 5 miles north on the base, and then another 5 miles on final to the runway threshold. Your current straight-line distance to the airport may be only 5 miles, but you have 15 flying miles to go.

I have two jets inbound from the west, straight-in to Runway 9. One is 10 miles out. The other is 20 miles out. Based on flying miles alone, the first jet has you beat: 10 flying miles versus your 15 flying miles. That jet is number one.

Speed comes into play considering the jet that's 20 miles out, 5 more flying miles than you. While the jet is currently doing 220 knots, it'll eventually slow to 140 knots on final approach. Assuming an average speed of 180 knots, it will arrive at the threshold in 6 minutes and 40 seconds. Your Cessna is indicating 100 knots. Even if you maintained 100 knots all the way to the runway threshold, it would still take you 9 minutes to fly 15 miles. The math doesn't lie. That second jet will likely be pulling into its parking spot while you're still on final.

How could I have made you number one? I could have cleared you for a visual approach, told you to keep your base turn as tight as possible and your speed up, and hope your definition of a tight base and keeping your speed up matches mine and the tower's. I could have slowed the jets down far earlier, perhaps issued them delay vectors, if necessary, to ensure I maintain my IFR separation from you on final.

All of that would increase my uncertainty, my workload, your workload, and the other pilots' workload, with little gain over just letting you and the other aircraft naturally fall into place. If you were in serious need—an emergency, serious weather rolling in, etc.—I'd make different choices.

Also, I'm not going to delay you indefinitely. On days where there's a never-ending string of fast-moving arrivals, I'll build a gap for you in my sequence. Mixing different aircraft types requires some patience on both sides.

Why does ATC issue vectors and how much thought goes into them?

I don't vector for funsies. Beyond ensuring safety, I view my job's main goal is to get aircraft where they want to go and do what they want to do, safely, with minimum delay. Whether you're IFR or VFR, on a filed routing or just fooling around, I will let you do your thing until there's a reason not to let you do that thing. Only then will I vector you.

What ATC Means When Asking You to Fly it Fast

Put simply, if I used the term "best forward speed" or "keep your speed up," give me the fastest speed you can comfortably fly. If you're concerned that's not enough, tell me quickly and succinctly.

Imagine you're flying a Beechcraft Baron G58 light twin and I'm working you into my sequence. I know the Baron's a pretty quick airplane, so I ask you, **"N123AB, maintain 190 knots or greater for sequence."** That aircraft's max cruise speed is 195 knots. It's a nice smooth day, so you push the throttles forward and read back, **"N123AB maintaining 190 knots or greater."**

However, what if you're encountering chop and turbulence? The Baron's maneuvering speed is 156 knots. How should you respond? **"Approach, N123AB, unable 190 knots due to turbulence. I can only maintain 155 knots."**

You can only do what you can do, and I can't ask you to do more. If what you can do is less than what I need to make my sequence or plan work, then I need to come up with a new plan. That's my job. It may involve vectors or other adjustments for you or other aircraft, but now I know what I'm working with.

Now I'm 100 percent guilty of telling aircraft in my tower pattern or landing sequence to **"keep your speed up"** in hopes of making a sequence work. When I request that, it's more about you not slowing unusually early. If you normally fly your downwinds at 100 knots, I don't want you prematurely slowing down to 80 or 90 knots. I'm realistic and familiar with the aircraft types I work and won't ask you for the moon and the stars.

In fact, I've been in the hot seat myself. Years back, I was doing my flight review in a Cessna 172 and was flying straight-in to the runway at the airport where I worked. My controller friend was working us in and told us to, **"Keep our speed up,"** due to jet traffic in trail. I was still intending to do an overall normal descent profile—what was comfortable for me—but just fly it a bit faster.

Instead, the instructor had me firewall the throttle. We held it high, no flaps, until the runway disappeared under the nose. Then he had me cut the power and do a diving forward slip to the runway. My butt was nearly off the seat as I stood on the rudder. As the runway rose up, I straightened out, dropped flaps, and landed. While a cool learning experience, yes, this was really pushing my comfort level at the time.

Here's what I suggest. Ask ATC for a specific speed rather than just "best forward"—and negotiate down if that's too fast. This way, the expectations are clear. That's true VFR or IFR.

I'm certainly appreciative when you and other pilots work with me and give me a little extra. However, I also know that, ultimately, no matter what I request, if the **"best forward speed"** you can give me is the standard speed you always use for whatever phase of flight you're in, then that is your best forward speed. Again: I want you safe and within your boundaries. I just need to know what those boundaries are.

And remember: All speed restrictions are deleted with an approach clearance (unless reissued with the clearance), and all automatically go away once you reach the FAF or 5 miles from the runway, whichever is closer. That's per **7110.65 5-7-1 (b)**. This ensures you'll have time to slow for a successful landing.

I can't just issue a turn and be done with it, though. I'm required to tell you why I'm vectoring you, and advise you of my post-vector plan for you. The **.65's 5-6-2 (b)** explains: *"When initiating a vector, advise the pilot of the purpose, and if appropriate, what to expect when radar navigational guidance is terminated."*

How many different ways can I finish the following sentence? **"N123AB, turn 20° left, vector for ..."**

Precipitation: I would much rather vector aircraft completely around observed or reported inclement weather than put them in a position where they need to fend for themselves. For instance, if I see there's an intense thunderstorm system 30 miles off your nose, I'll give you vectors clear of it. A 10° vector early can save a whole lot of vectors and precipitation calls later. Of course, sometimes there's no option but to hunt and peck your way through the weather. If a bunch of scattered showers are bracketing your destination, I may be stuck vectoring you toward gaps or weaker spots in the precipitation.

Airspace: This can be a catch-all for a number of things:

My adjacent radar facilities may have LOAs specifying how aircraft must be handed off on headings rather than their filed routings. Once the next facility takes control, then they'll put the aircraft back on their routing.

I may need to steer you around Special Use Airspace, like an active MOA, active Restricted Areas, or a TFR.

If you're getting close to a neighboring radar sector or tower's airspace and they can't accept you as a point out or handoff (or I'm too busy to call right then), I'll vector you away from their airspace.

Noise Abatement: Arrivals or departures may require vectors to comply with noise abatement procedures established in LOAs or SOPs.

Traffic: Easy one here. If I observe that aircraft are currently, or will potentially be, in conflict, vectors for one or both aircraft can resolve the matter. I will also vector aircraft away from known airborne threats, like a parachute jump aircraft that's about to unload its human meat missiles over its jump zone.

Sequence: When I'm working arrivals, I'm typically getting a mix of aircraft either cleared direct to their destination airport, inbound on a STAR procedure, or VFRs that may need to get meshed with my IFR traffic. All of these aircraft may be coming from different directions. I'll use vectors and speed control to sequence them to each appropriate arrival runway. Even if there are no other aircraft but you, a tower's LOAs may require vectors. For instance, the towers within my approach control's airspace require that all IFR jets be vectored to at least a 5-mile final. A jet will get those vectors even if it's the only aircraft in the sky.

In-Trail Spacing: Departures and enroute aircraft need to get spaced out and lined up, too. For instance, if I get simultaneous departures off separate airports, all going out the same piece of airspace (what ATC calls an "airspace gate"), I must organize them into a line to properly feed them to the next controller. Or maybe a center sector needs 20 miles between all aircraft entering its airspace. It can take some vectors to achieve that.

Pilot Request: Pilots can request a vector from ATC at any time. Perhaps you just received an amended clearance while airborne and are struggling to get it loaded into your GPS. **"Approach, N123AB, request vector to [first fix]. I'll go direct when able."** Maybe you're VFR, trying to find your destination airport, but you're looking straight into the setting sun. The airport's likely depicted on ATC's radar, so the controller can give you a vector to it.

Controllers aren't issuing vectors for amusement. There's a lot of thought, problem-solving, and regulatory compliance going on behind the scenes.

How does ATC get me back on course?

Well, I've vectored you so much your GPS track looks like a Winter Olympics slalom course. At last, whatever traffic, weather, or obstructions I vectored you around are no longer an impediment. I'm ready to hand the reins back to you. This can be phrased several ways, depending on whether you're VFR, IFR, or IFR on a SID or STAR.

VFR: There's not much to consider if you're VFR. You aren't tied to a filed routing and I therefore don't need to connect you back to a known point or airway. When I say, **"N123AB, resume own navigation,"** you're free to turn as you wish. Don't forget, though: Any existing altitude restrictions remain in effect.

IFR: When you're IFR, however, **"resume own navigation"** is ambiguous. There are many ways you could rejoin your routing. We don't want surprise navigation under IFR. So to avoid you having questions, I'll use one of the following:

Cleared Direct [fix]. I'll clear you direct to a point on your routing, be it a NAVAID, waypoint, fix, etc. This is my preferred method, as it's quick and concise. Here I'd say, **"N123AB, cleared direct DELCO."** You go direct DELCO, and then direct to KHUF as that's the next point in your routing.

Fly heading, join [airway] [direction]: I can give you vectors to (re)join an airway in your routing. For instance, **"Fly heading 100, join T305 eastbound."** You'll intercept T305 and carry on your merry way.

Fly heading [heading], when able to proceed direct [fix]: Sometimes I'll give you a heading to get you started in the right direction, capping it off with a direct-to-fix instruction. **"N123AB, fly heading 090. When able, proceed direct DELCO."** I'll often use this when there's potential traffic or another concern, and I want to deconflict you quickly, without waiting for you to program your GPS appropriately. This is the common phraseology for vectors around weather as well.

SIDs/STARs: These departure and arrival procedures have additional layers to consider. They're already an often-complex latticework of waypoints, altitudes, crossing restrictions, and speeds. Getting you back on the rails after a vector takes some extra planning and thought.

> **TIP** You'll often hear controllers say, **"Proceed on course."** I've been guilty of it myself at times. As ubiquitous as it is, that isn't correct phraseology. The proper catch-all phraseology for cutting you loose on your own recognizance is: **"Resume own navigation."**

If I intend for you to rejoin the procedure at some point, I need to tell you so when I initiate the vector per **7110.65 5-6-2 (d)** on vectoring methods, : *"When vectoring or approving an aircraft to deviate off of a procedure, advise the pilot if you intend on clearing the aircraft to resume the procedure."*

Imagine you were climbing via the (imaginary) FIXER 2 Departure that stretches for a hundred miles' worth of waypoints and altitudes. Its first three waypoints are FIXAA, FIXBB, and FIXCC. There's a small storm hovering between FIXAA and FIXBB. **"N123AB,"** I say, **"fly heading 310, vector for precipitation. Expect to rejoin the SID at FIXCC."** That makes it clear that you shouldn't dump the whole SID out of your system so that you're ready to fly direct to FIXCC later.

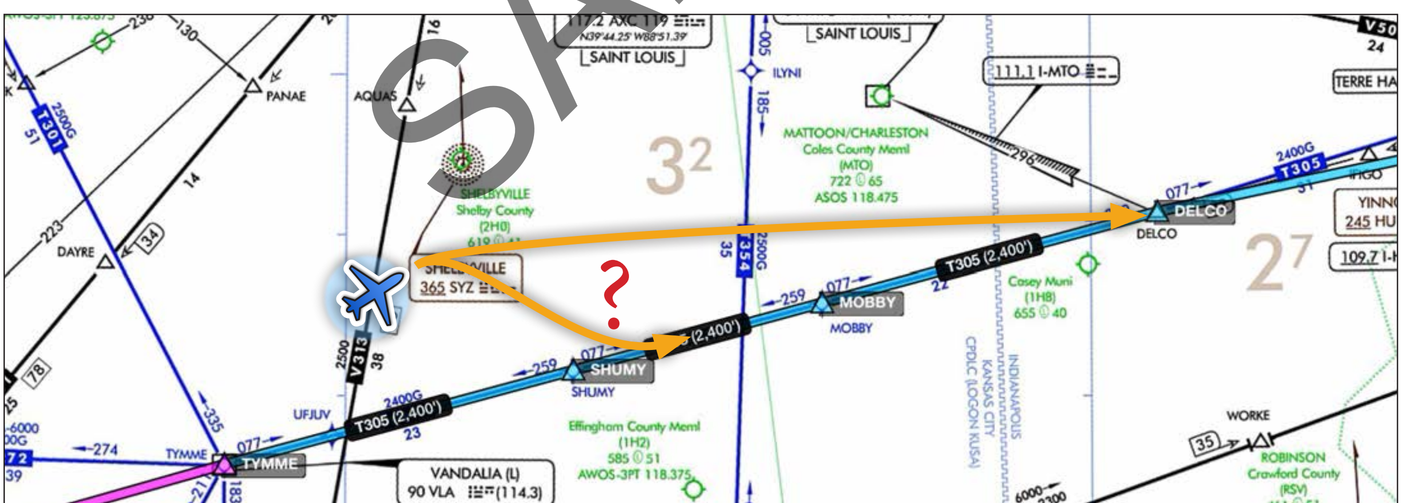

Imagine you're filed from St. Louis Regional Airport (KALN) in Missouri to Terre Haute, IN (KHUF) via: "KALN TYMME T305 DELCO KHUF." Somewhere after you passed TYMME, I vectored you north of the airway for weather. Now I tell you, "N123AB, resume own navigation." *Do you: A) Turn right to rejoin T305? B) Proceed direct DELCO? C) Something else? What if terrain or obstructions are potential factors, depending on how you proceed? Here's where I remind you of* ***91.123****: "When a pilot is uncertain of an ATC clearance, that pilot shall immediately request clarification from ATC." Not sure? Ask.* **"Approach, N123AB, did you want me to go direct DELCO or rejoin T305?"**

What about altitude and speed restrictions published in the procedure? The **.65 5-6-2 (d)** continues: *"After a climb via or descend via clearance has been issued, a vector/deviation off of a SID/STAR cancels all published altitude and speed restrictions on the procedure."* The avionics on some aircraft may be incapable of processing crossing altitude restrictions once the aircraft is pulled off the SID/STAR midstream, defaulting to the procedure's published top or bottom altitude.

How does this get fixed? The **.65** section **5-6-2 (f)** states: *"Aircraft instructed to resume a procedure which contains published crossing restrictions (SID/STAR) must be issued/reissued all applicable restrictions or be instructed to Climb Via/Descend Via."* So to tie you back into your procedure, I say, **"N123AB, cleared direct FIXCC, cross FIXCC at or above 9000, then climb via the FIXER 2 Departure."**

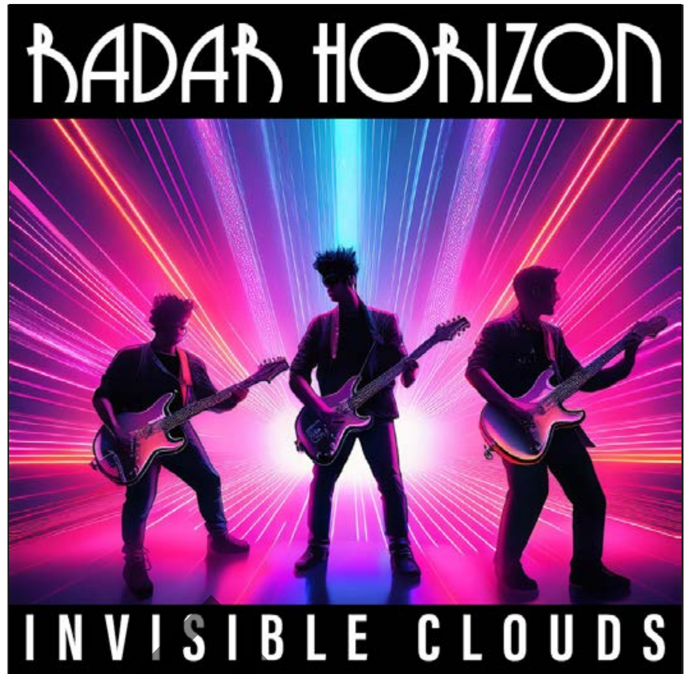

Radar Horizon sounds like a progressive rock band name, doesn't it?

ATC keeps talking about the MVA. What is that?

The Minimum Vectoring Altitude (MVA) is the lowest altitude that ATC can vector IFR aircraft (with some exceptions). It's generally defined by two things: ***distance*** from our radar/sensor site and the ***height*** of obstructions/terrain within our airspace.

The farther from a radar you get, the more its effectiveness at low altitude gets impacted by ground clutter, atmospheric conditions, and the curvature of the Earth. The dividing line between where the radar can still see and where things are hidden from it is called the radar horizon. The MVA will be set above that horizon to guarantee radar detection.

When there are no obstructions or prominent terrain, the MVA in non-mountainous terrain is 1000 feet above the highest obstacle, which tends to be about 1500 to 2000 feet AGL. Add 1000 feet for mountainous areas. For instance, Blue Grass Airport in Lexington, KY (KLEX) has a field elevation of 979 feet MSL. The MVA for its surrounding airspace is 2500 feet MSL. That keeps aircraft 1500 feet above all population areas and above the radar horizon.

Now obstructions and terrain have entered the chat. Here's a tip for the former: Take out your chart, point to an obstruction on it (like an antenna), and figure out its height above sea level. Now add 1000 feet. That's likely the MVA in that area. For standalone obstructions, the MVA "bubble" extends out to a radius of several miles and 1000 feet above the object.

Mountainous terrain gets an added bump. Rather than 1000 feet, the MVA is 2000 feet above the highest peak in the area. A little extra bit of protection and coverage never hurt anyone.

As I said earlier, there are some exceptions where we can vector below the MVA. First, VFR aircraft are less restricted by MVAs. I can vector them below the MVA as long as I don't simultaneously assign them an altitude that's below the MVA. If I must issue an altitude under the MVA, then I need to cancel the vectors and have the VFR aircraft resume its own navigation. It's part of the see-and-avoid aspect. VFR traffic below the MVA must be free to avoid clouds and obstructions either vertically or laterally. If one is restricted, they must retain the other option.

I can also vector departing IFR aircraft below an MVA if certain conditions are met. Example? If I have assigned you an altitude at least 1000 feet above a relevant obstruction, then I can vector you to maintain or achieve at least 3 miles of separation from that obstruction. I must depict these obstructions on my scope, of course. Once you're at or above the MVA, though, I can't vector you below it again.

We can get really into the weeds with additional exceptions, such as go-around headings, SIDs, and Diverse Vector Areas (DVAs). The **.65's (5-6-3)** Vectors Below Minimum Altitude lays out a number of examples. The short version: We can vector IFR departures if the FAA has reviewed the surrounding area, obstructions, and terrain, and then authorized procedures and headings that ensure appropriate separation.

MVA and MIA charts for many ATC facilities are publicly available online. While there is overlap in design between the two, there are numerous variances. Some EFB apps can import MVA and MIA data overlaid on the Sectional or Enroute Chart. This is much more practical than a cryptic map that only makes sense to someone working that facility.

Since we're talking about minimum altitudes, I should note that MVAs are used by terminal facilities (towers and approach controls). Centers instead use a similar concept called *Minimum IFR Altitude (MIA).* Centers oversee vast amounts of airspace and may have gaps in their radar coverage. While there is a lot of overlap in MIA and MVA altitudes, MIA charts are designed to be used without radar (hence the lack of "V-for-vectoring" in the acronym).